Tropical Geometry

by

David E Speyer

A.B. (Harvard University) 2002

A dissertation submitted in partial satisfaction of the
requirements for the degree of
Doctor of Philosophy

in

Mathematics

in the

GRADUATE DIVISION
of the
UNIVERSITY of CALIFORNIA, BERKELEY

Committee in charge:

Professor Bernd Sturmfels, Chair
Professor Allen Knutson
Professor Steven Evans

Spring 2005

The dissertation of David E Speyer is approved:

Chair Date

Date

Date

University of California, Berkeley

Spring 2005

Tropical Geometry

Abstract

Tropical Geometry

by

David E Speyer

Doctor of Philosophy in Mathematics

University of California, Berkeley

Professor Bernd Sturmfels, Chair

Let K be an algebraically closed field complete with respect to a nonarchimedean valuation $v : K^* \to \mathbb{R}$. The reader should think of K as the field of Puiseux series, $\bigcup_{n=1}^{\infty} \mathbb{C}((t^{1/n}))$ and v as the map that assigns to a power series the exponent of its lowest degree term. Let X be a subvariety of the torus $(K^*)^n$. Then $v(X)$ is (essentially) a polyhedral complex which we term the tropicalization of X and its combinatorial properties reflect the geometry of X.

In this thesis, we first discuss structural results about the tropicalization of X for an arbitrary X. We then turn to the special cases of linear spaces, Grassmannians and curves; in each case trying to describe as explicitly as possible the possible combinatorics of the tropicalization. In the case of linear spaces, our investigations lead us to a combinatorial theory of tropical linear spaces. This theory leads to interesting combinatorial problems about matroids. In the case of curves, the relevant combinatorial objects are "zero tension curves". We prove a few combinatorial results but concentrate on the other problem of determining which zero tension curves actually occur as tropicalizations.

Professor Bernd Sturmfels
Dissertation Committee Chair

Contents

Acknowledgements

Throughout my study of this material, I have benefitted from the willingness of many mathematicians to discuss their work and that of others with me. I will doubtless omit some important names, but Paul Hacking, Allen Knutson, Grisha Mikhalkin, Ezra Miller, Arthur Ogus and my advisor Bernd Sturmfels stand out particularly. I also must thank Bernd for his guidance in editing and planning this entire project. Finally, my thanks to Erin Larkspur, who has kept me sane throughout the process.

Chapter 1

Introduction

The fundamental observation of tropical geometry is that the solution sets to algebraic equations, when plotted in logarithmic coordinates, look roughly like polyhedral complexes. This behavior was first made precise in Bergman's paper [3] and spelled out more intuitively in Viro's paper [46]. The aim of tropical geometry is to make this analogy into a precise and useful tool.

In order to do this, we switch from considering the logarithm defined on the real or complex numbers to considering a valuation on a non-archimedean field (defined within). Now the image of the valuation will not only resemble a polyhedral complex but will actually be one. If X denotes the solution space of our family of equations then the image under the valuation map is called the tropicalization of X and denoted $\operatorname{Trop} X$.

In the next chapter we give several descriptions of this set, prove that it is indeed polyhedral, prove some of its properties and introduce a degeneration of the solution set of the original equations whose components are indexed by the faces of the polyhedral complex.

We then turn to applications of our theory to the fundamental examples of algebraic geometry: linear spaces, Grassmannians and curves. In chapter 3, we describe work of B. Sturmfels and myself to try to compute the tropicalization of the Grassmannian $G_{d,n}$ – the space of d planes in n space. We have a complete and elegant description in the case $d = 2$ and we also have a complete description of the case $(d, n) = (3, 6)$ by direct computation.

It turns out that, just as in the classical case, the tropicalization of the Grassmannian parameterizes the possible tropicalizations of d planes in n space and that this is the best method for describing results about $\operatorname{Trop} G_{d,n}$. We have also found that there is a notion of "tropical linear space", which is more general than actual tropicalizations of linear spaces and which is purely combinatorial. In chapter 4 we first introduce this notion, prove several equivalent characterizations of it, and introduce combinatorial analogues of intersection and dualization. We then discuss how actual tropicalizations of linear spaces occur among all the tropical linear spaces and show that $\operatorname{Trop} G_{d,n}$ precisely parameterizes the collection of tropicalizations of linear spaces.

Next, we then turn towards investigating the combinatorics of tropical linear spaces in their own right. This section should be appealing to readers who prefer combinatorics to algebraic geometry. Our main target is to prove the "f-vector conjecture", which describes explicitly how complicated the tropicalization of a d-plane in n-space can be. While we have not managed to prove the conjecture in its full generality, we prove many special cases, including proving it for all tropical linear spaces built out of a starting collection of hyperplanes by repeated intersection and dualization, proving several particular cases of the bounds and exhibiting a particularly elegant class of tropical linear spaces, the "tree spaces", which achieve these bounds.

In Chapter 5, we turn to the study of the tropicalizations of curves. This field was pioneered by Mikhalkin and is to a large extent responsible for the recent popularity of tropical thinking. We take a very different approach than Mikhalkin and, rather than study curves *via* deformation of complex structure and other symplectic techniques, we instead use Tate and Mumford's theories of non-archimedean uniformization. Our main result, which has been conjectured by Mikhalkin and will appear in a future paper of his proven by quite different methods, is that any graph in $\mathbb{R}^n$ which satisfies the local condition on the graph known as the zero tension condition and a global condition called being "ordinary", which amounts to a certain matrix having full rank, actually occurs as the tropicalization of a curve in the n-dimensional torus.

Chapter 2

General Constructions

2.1 Definitions of the Tropicalization of a Variety

Let K be an algebraically closed field complete with respect to a non-trivial valuation $v : K \to \mathbb{R} \cup \{\infty\}$. Recall that "non-trivial valuation" means that

$$\begin{aligned} v(x+y) &\geq \min(v(x), v(y)) \\ v(xy) &= v(x) + v(y) \\ v(0) &= \infty \\ v(1) &= 0 \\ v(x) & \quad \text{takes values other than 0 and } \infty \end{aligned}$$

Note that the second and fourth property automatically imply that, if $x \neq 0$ then $v(x) + v(x^{-1}) = 0$ so $v(x) \neq \infty$. When speaking about valuations from a ring which is not a field, as we will have occasion to do in Theorem 2.1.2, we permit

$v(x) = \infty$ for nonzero x.

That K is complete with respect to v means that, if $\alpha_1,\ \alpha_2\ \ldots$ is a seqeunce of elements of K with $\lim_{i,j\to\infty} v(\alpha_i - \alpha_j) = \infty$ then there is an element β of K with $\lim_{i\to\infty} v(\beta - \alpha_i) = \infty$.

In order to maintain compatibility with the various valued fields that occur in mathematics, we have not assumed v to be surjective. This will cause many technical frustrations. Our approach is to be scrupulous about handling such issues in this chapter but to occasionally assume that relevant values are in the image of v in later sections. This issue never has any deep effect.

We set the following notations: $\mathcal{R}$ will be the local ring $v^{-1}(\mathbb{R}_{\geq 0} \cup \{\infty\})$, $\mathcal{M}$ the maximal ideal $v^{-1}(\mathbb{R}_{>0} \cup \{\infty\})$ of $\mathcal{R}$ and κ the field $\mathcal{R}/\mathcal{M}$. For most purposes, it is best to think of K as the field $\bigcup_{n=1}^{\infty} \mathbb{C}((t^{1/n}))$ of Pusieux series, or as the reader's favorite algebraically closed field of power series. It will be convenient, however, to have the flexibility to talk about an arbitrary K.

Since K is algebraically closed, it is closed under the taking of n^{th} roots for all $n \in \mathbb{Z}$. Thus, K^* and $v(K^*)$ are divisible groups and we can find a section of the surjective map of groups $K^* \to v(K^*)$. We fix such a section $w \mapsto t^w$ from $v(K^*) \to K^*$. Explicitly, this means that $t^{w+w'} = t^w t^{w'}$ and $v(t^w) = w$. Once again, the reader should think of the case where K is a field of power series, so that the notation t^w can simply be thought of as a monomial in K.

One notion that we will need repeatedly is the notion of an initial ideal.

Let Y be a toric variety over K with dense torus $(K^*)^n$ (we will usually be thinking about the torus itself and almost always about the torus, affine space or projective space), X a closed subscheme of Y and $w \in v(K^*)^n$. Let $\mathcal{Y}$ be the toric variety over $\mathcal{R}$ associated to the same fan and let $\mathcal{X}$ be the closure of $t^w \cdot X \subset (K^*)^n$ in $\mathcal{Y}$. We define $\operatorname{in}_w X = \mathcal{X} \times_{\mathcal{R}} \kappa$.

While $\operatorname{in}_w X$ depends on the choice of $t^\cdot$, it does so only in a trivial manner: a different choice of $t^\bullet$ amounts to translating $\operatorname{in}_w X$ by an element of $(\kappa^*)^n$. Similar comments will apply to most of our constructions.

Let $f = \sum_{a \in A} f_a x^a$ be a nonzero Laurent polynomial in $K[x_1^{\pm 1}, \dots, x_n^{\pm 1}]$ and let $w \in v(K^*)^n$. Let $W = \min_{a \in A} (\sum w_i a_i + v(f_a))$, so the polynomial $t^{-W} f(t^{w_1} x_1, \dots, t^{w_n} x_n)$ is in $\mathcal{R}[x^{\pm 1}]$ and has nonzero image in $\kappa[x^{\pm 1}]$, let $\operatorname{in}_w f$ be this image. Let X be a subvariety of the torus $(K^*)^n$ defined by the ideal $I \subset K[x_1^{\pm 1}, \dots, x_n^{\pm 1}]$.

Proposition 2.1.1. *The ideal of* $\operatorname{in}_w X$ *is spanned over* κ *by the polynomials* $\operatorname{in}_w f$ *as* f *ranges over* $I \setminus \{0\}$.

Proof. The definition of closure tells us that the ideal of $\operatorname{in}_w X$ given by the ideal $(I \cap \mathcal{R}[x_1^{\pm}, \dots, x_n^{\pm}]) \otimes_{\mathcal{R}} \kappa$ in $\kappa[x_1^{\pm 1}, \dots, x_n^{\pm 1}]$. Let f and W be as in the preceding paragraph with $f \in I \setminus \{0\}$; we will write $W(f)$ when necessary. We have $f \in \mathcal{R}[x^{\pm}]$ if and only if $W \geq 0$. If $f \in \mathcal{R}[x^{\pm}]$ then its image in $(I \cap \mathcal{R}[x_1^{\pm}, \dots, x_n^{\pm}]) \otimes_{\mathcal{R}} \kappa$ is 0 if $W > 0$ and is $\operatorname{in}_w f$ if $W = 0$. So the ideal of $\operatorname{in}_w f$ is by definition spanned over κ by the polynomials $\operatorname{in}_w f$ for which $f \in I$ and $W(f) = 0$. But, for any

$f \in I \setminus \{0\}$, $W(t^{-W(f)}f) = 0$, $t^{-W(f)}f \in I$ and $\mathrm{in}_w f = \mathrm{in}_w\left(t^{-W(f)}f\right)$. So the set of polynomials of the form $\mathrm{in}_w f$ with $f \in I \setminus \{0\}$ is the same as the subset of polynomials of the form $\mathrm{in}_w f$ with $f \in I \setminus \{0\}$ and $W(f) = 0$. $\square$

We use the above proposition to extend the definition of $\mathrm{in}_w X$ to the case where $w \notin v(K^*)$.

Let $X \subset (K^*)^n$ be a closed subvariety of the torus with $I(X) \subset K[x_1^{\pm 1}, \ldots, x_n^{\pm 1}]$ the corresponding ideal. In the remainder of this section, we will define four subsets of $v(K^*)^n$ and prove that they are equal. We will then define $\mathrm{Trop}\, X$ to be the closure of these sets in $\mathbb{R}^n$.

Theorem 2.1.2. *The following subsets of $v(K^*)^n$ are equal:*

1. *The set of all $(v(u_1), \ldots, v(u_n))$, where $(u_1, \ldots, u_n) \in X(K)$.*
2. *The set of all points of $v(K^*)^n$ of the form $(\tilde{v}(x_1), \ldots, \tilde{v}(x_n))$ where $\tilde{v} : \mathcal{O}(X) \to \mathbb{R} \cup \{\infty\}$ is a valuation extending v.*
3. *The set of all $w \in v(K^*)^n$ such $\mathrm{in}_w f$ is not a monomial for any $f \in I(X) \setminus \{0\}$.*
4. *The set of all $w \in v(K^*)^n$ such that $\mathrm{in}_w X \neq \emptyset$.*

Proof. We provisionally term these sets T_1, T_2, T_3 and T_4 and proceed to show that $T_1 \subseteq T_2 \subseteq T_3 \subseteq T_4 \subseteq T_1$.

$T_1 \subseteq T_2$: Define $\tilde{v}$ by $\tilde{v}(f) = v(f(u_1, \ldots, u_n))$ for every $f \in \mathcal{O}(X) = K[x^{\pm}]/I(X)$.

$T_2 \subseteq T_3$: Suppose that $w \in T_2$. Let $f = \sum_{a \in A} f_a x^a \in I(X) \setminus 0$, let $\tilde{v}$ be a valuation $\mathcal{O}(X) \to \mathbb{R} \cup \{\infty\}$ extending v with $w_i = v(x_i)$. In $\mathcal{O}(X)$, we have $f(x_1, \ldots, x_n) = 0$ which means that the minimum of the numbers $\tilde{v}\left(f_a x_1^{a_1} \cdots x_n^{a_n}\right)$ is not unique. But

$$\tilde{v}(f_a x_1^{a_1} \cdots x_n^{a_n}) = \tilde{v}(f_a) + \sum a_i \tilde{v}(x_i) = v(f_a) + \sum a_i w_i,$$

so saying that the minimum of this quantity as a ranges over A is not unique is exactly saying that $\mathrm{in}_w f$ is not a monomial.

$T_3 \subseteq T_4$: Suppose that $w \notin T_4$. By the Nullstellensatz, to say that $\mathrm{in}_w X \neq \emptyset$ is to claim that $\mathrm{in}_w I(X) \not\ni 1$. Suppose that $1 \in \mathrm{in}_w I(X)$, and write $1 = \sum_{i=1}^r \mathrm{in}_w f^i$ for $f^i = \sum_{a \in A_i} f_a^i x^a \in I(X)$. For each i, set $W_i = \min_{a \in A_i} (v(f_a) + \sum a_i w_i)$, so $F := \sum_{i=1}^r t^{-W_i} f^i \equiv 1 \mod \mathcal{M}$. Then $\mathrm{in}_w F = 1$, so $w \notin T_3$.

$T_4 \subseteq T_1$: This is the only difficult part. Recall our definition of $\mathcal{X}$ as the closure in the torus over $\mathcal{R}$ of $t^w X \subset (K^*)^n$, and $\mathrm{in}_w X$ as the fiber of $\mathcal{X}$ over κ. We have assumed that $\mathrm{in}_w X \neq \emptyset$, let $(\overline{u}_1, \ldots, \overline{u}_n) \in \mathrm{in}_w X$. We must prove there is a point $(x_1, \ldots, x_n) \in X$ with $v(x_i) = w_i$. We actually prove the following stronger result:

Lemma 2.1.3. *If $w \in v(K^*)^n$ and $(\overline{u}_1, \ldots, \overline{u}_n)$ is a point of $\mathrm{in}_w X$ then there is a point $(u_1, \ldots, u_n) \in X$ with $u_i = \overline{u}_i t^{w_i}$ + higher order terms.*

This result resembles Hensel's lemma, but allows us to treat more general varieties at the expense of treating less general fields. Hensel's lemma has as an additional hypothesis that the fiber over κ is smooth and can therefore avoid assuming K to be algebraically closed.

Proof. By multiplying X by t^{-w}, we may assume that $w = 0$. Let $\overline{m}$ be the maximal ideal

$$\overline{m} = \langle x_1 - \overline{u}_1, \ldots, x_n - \overline{u}_n \rangle + \mathcal{M}$$

in $\mathcal{R}[x^{\pm}]$. Consider the localization $\mathcal{O}(\mathcal{X})_{\overline{m}}$ of $\mathcal{O}(\mathcal{X})$ at $\overline{m}$; this is $S^{-1}\mathcal{O}(\mathcal{X})$ where $S = \{f \in \mathcal{R}[x^{\pm}] : f \not\equiv 0 \mod \overline{m}\}$. Since $\mathcal{X}$ is defined as the closure of its generic fiber over the one dimensional scheme $\operatorname{Spec}\mathcal{R}$ we have that $\mathcal{O}(\mathcal{X})$ is flat over $\mathcal{R}$ and the localization $\mathcal{O}(\mathcal{X})_m$ is as well. Therefore, every minimal prime of $\mathcal{O}(\mathcal{X})_m$ lies over K. Let P be such a prime. Then there is a maximal ideal M of $\mathcal{O}(X) = K[x^{\pm 1}]/I$ containing $P \otimes_{\mathcal{R}} K$; as K is algebraically closed, this ideal is of the form

$$\langle x_1 - u_1, \ldots, x_n - u_n \rangle$$

for some $u \in X$. We must show that $u_i \in \mathcal{R}$ and $u_i \equiv \bar{u}_i \mod \mathcal{M}$.

If $u_i \notin \mathcal{R}$ then $u_i^{-1} \in \mathcal{M}$ and $u_i^{-1}x_i - 1 \equiv -1 \mod \mathcal{M}$ so $u_i^{-1}x_i - 1 \in S$. As the elements of S are units in $\mathcal{O}(\mathcal{X})_m$, this contradicts that $u_i^{-1}x_i - 1$ is in a maximal ideal of $\mathcal{O}(\mathcal{X})_m \otimes_{\mathcal{R}} K$. Similarly, if $u_i \in \mathcal{R}$ but $u_i \not\equiv \overline{u}_i \mod \mathcal{M}$ then $x_i - u_i \in S$ and the same contradiction applies. □

□

We define $\operatorname{Trop} X$ to be the closure of the set defined in the above theorem in $\mathbb{R}^n$. If we are given X in a toric variety over K, we write $\operatorname{Trop} X$ as shorthand for $\operatorname{Trop}(X \cap (K^*)^n)$. $\operatorname{Trop} X$ seems to have first been defined in Bergman's paper [3]. Bergman deals with the constant coefficient case and first gives a definition in terms of the amoeba of X, which he then relates to a modification of definition T_2 above. Bieri and Groves, in [6], used essentially definition T_2 over a general valued field.

Our order of definitions: first defining a subset of $v(K^*)^n$ and then defining $\operatorname{Trop} X$ as its closure in $\mathbb{R}^n$ leaves the possibility that there are points of $\operatorname{Trop} X \cap v(K^*)^n$ which do not meet the conditions of the above theorem. The proposition below shows that there are no ambiguities of this sort.

Proposition 2.1.4. *The set described in the above theorem is closed in $v(K^*)^n$. Moreover,* $\operatorname{Trop} X$ *can be described as the set of $w \in \mathbb{R}^n$ such that*

1. $w = (\tilde{v}(x_1), \ldots, \tilde{v}(x_n))$ *where* $\tilde{v} : \mathcal{O}(X) \to \mathbb{R} \cup \{\infty\}$ *is a valuation extending* v*.*

2. $\operatorname{in}_w f$ *is not a monomial for any* $f \in I(X) \setminus \{0\}$*.*

3. $\operatorname{in}_w X \neq \emptyset$*.*

Proof. The set of w such that $\operatorname{in}_w f$ *is* a monomial is the union over the finitely many monomials in f of the set of w for which that particular monomial is the leading term of f. The condition that a particular monomial is the leading term is open, thus, the condition that $\operatorname{in}_w f$ is *not* a monomial is closed. So, from

Definition 3 of Theorem 2.1.2, we see that the set defined in the above theorem is closed in $v(K^*)^n$ and that (2) above characterizes $\operatorname{Trop} X$.

The proof that $T_3 \subseteq T_4$ above also shows that (2) implies (3). It is even easier to show that (3) implies (2): If $\operatorname{in}_w f$ is a monomial for some $f \in I$ then, since monomials are units in $\kappa[x^{\pm}]$, the ideal of $\operatorname{in}_w X$ is (1) and $\operatorname{in}_w X$ is empty.

Finally, suppose that $w \in \operatorname{Trop} X$. We can extend K to a field L with valuation such that $w \in v(L^*)^n$. (Proof: make a degree n transcendental extension $K(u_1, \ldots, u_n)$ of K and set $v(\sum a_I u^I) = \min_I(v(a_I) + < u, w >)$ where $a_I \in K$. It is not hard to check that this is multiplicative, and hence has a well defined extension to the fraction field, and that the extended v is still a valuation.) Replacing I with $I \otimes_K L$ will not change the truth of condition 3 in the previous theorem, so $w_i = x_i$ for some $(x_1, \ldots, x_n) \in X(L)$. Then we can define $\tilde{v} : \mathcal{O}(X) \to \mathbb{R}$ by $f \to v(f(x_1, \ldots, x_n))$. □

The characterization of $\operatorname{Trop} X$ in terms of the set of all valuations extending v will not be used in the future. I include it because it shows

Proposition 2.1.5. $\operatorname{Trop} X$ *is a continuous surjective image of the analytic space associated to* X *in the sense of Berkovich [4].*

Proof. The points of this space are, by definition, the norms on X extending the norm $e^{-v(\cdot)}$ on K. There is a bijection between valuations and norms by $\tilde{v} \to e^{-\tilde{v}(\cdot)}$ where we take $e^{-\infty} = 0$. The functions $e^{-\tilde{v}(\cdot)} \to \tilde{v}(x_i)$, for $1 \leq i \leq n$, are continuous. □

This observation (using the language of Tate's rigid analytic spaces rather than Berkovich's analytic spaces) is used in [17]. I think that there is a good opportunity for cross fertilization between the tropical and analytic geometry communities, as tropical geometry has not yet developed the sort of abstract geometrical theory the analytic scholars have and, as far as I have found, the analytic community is not utilizing the polyhedral nature of $\operatorname{Trop} X$.

2.2 Polyhedral Structure of $\operatorname{Trop} X$

Let Y be the torus $(K^*)^n$, the affine space K^n or the projective space $\mathbb{P}^n_K$ and X a closed subscheme of Y. We will say that a polyhedral subdivision Σ of $\mathbb{R}^n$ is adapted to X if, whenever w and w' lie in the relative interior of the same face of Σ, we have $\operatorname{in}_w X = \operatorname{in}_{w'} X$. The main result of this section is

Theorem 2.2.1. *There is a polyhedral subdivision Σ of $\mathbb{R}^n$ which is* adapted to *X. We may choose that each face of Σ be defined by inequalities of the form $\sum_{i=1}^n a_i w_i \geq c$ with $a_i \in \mathbb{Z}$ and $c \in v(K^*)$.*

This result essentially extends the results of [27]; our method of argument is essentially that of [2]. These papers basically treat the case that K is a power series field over κ and X is defined over κ. We term this case the *constant coefficient* case – explicitly, we will say that we are in the constant coefficient case when κ embeds into K and X is defined over K. In the constant coefficient case the construction that follows will show that we can take Σ, and thus $\operatorname{Trop} X$, to be a fan. In this case, Σ is called the Gröbner fan of X. See also [41] for more on computing this fan.

Proof. Suppose that Y' is a toric variety that contains Y as an open dense subvariety and let X' be the closure of X in Y'. Then $\operatorname{in}_w X = (\operatorname{in}_w X') \cap Y$ so proving the theorem for Y' proves it for Y. Thus, we may assume that Y is projective space. Write $Y = \operatorname{Proj} K[y_0, \ldots, y_n]$ and $X = \operatorname{Proj} K[y]/I$ for some homogeneous ideal I.

Let $w \in v(K^*)^n$. We define the ideal $\operatorname{in}_w I$ to be the homogeneous ideal $\operatorname{in}_{\tilde{w}} I$ where $\tilde{w}$ is an arbitrary preimage of w under the map $\mathbb{R}^{n+1} \to \mathbb{R}^{n+1}/(1, \ldots, 1)$. Note that this is a *not necessarily saturated* homogeneous ideal, whose saturation is the homogeneous ideal of $\operatorname{in}_w X$.

As w varies, all of the ideals $\operatorname{in}_w I$ have the same Hilbert function $h : \mathbb{Z}_{\geq 0} \to \mathbb{Z}$ given by $a \mapsto \dim(K[y]/\operatorname{in}_w I)_a$. Given a fixed Hilbert function, there is a set of degrees $a_1, \ldots, a_k$ such that any homogeneous ideal I with that Hilbert function is generated in degrees $a_1, \ldots, a_k$. (Proof: it is enough to prove this for monomial ideal I, as every homogenous ideal I has a Gröbner degeneration to a monomial ideal with the same Hilbert function and generators in the same degree, see, for example, chapter XVI of [18]. There are only finitely many monomial ideals with a given Hilbert function, see corollary 2.2 of [23].) Thus, as I varies, $\operatorname{in}_w I$ is determined by $(\operatorname{in}_w I)_{a_i}$ where $1 \leq i \leq k$. Our theorem will follow from the lemma below:

Lemma 2.2.2. *Let $V = K^M$ be a finite dimensional K vector space of dimension M equipped with an action of $(K^*)^n$; choose an eigenvector decomposition $V = \bigoplus_{i=1}^M Kv_i$ with v_i an eigenvector of character χ_i for this action. For U a subspace of V, let $U_0 \subset \kappa^M$ be defined by $(U \cap \bigoplus \mathcal{R}v_i) \otimes_{\mathcal{R}} \kappa$. As w varies through $v(K^*)^n$, the vector space $(t^w U)_0$ only takes on finitely many values and, partitioning $\mathbb{R}^n$ into equivalence classes by the value of $(t^w U)_0$, these classes are the relatives interiors of the faces of a polyhedral complex.*

This lemma proves our result by taking $V = K[y]_{a_i}$ and $U = I_{a_i}$ and then taking the simultaneous refinement of the resulting polyhedral complexes.

Proof. Write D for the dimension of U. Note that the Plücker coordinates of U_0 in the basis v_i are given by $p_{i_1 \cdots i_D}(U_0) = 0$ if $v(p_{i_1 \cdots i_D}(U))$ is not minimal among all valuations of Plücker coordinates of U and, $p_{i_1 \cdots i_D}(U_0) = p$ where $p_{i_1 \cdots i_D}(U) = t^W p + \cdots$ for some W if $p_{i_1 \cdots i_D}(U)$ is one of these Plücker coordinates of minimal valuation. The Plücker coordinates of $t^w \cdot U$ are given by $p_{i_1 \cdots i_D}(t^w \cdot U) = t^{\sum_{r=1}^D <\chi_i, w>} p_{i_1 \cdots i_D}(U)$. We thus see that the Plücker coordinates of $(t^w \cdot U)_0$ depend only on which of the valuations $v(p_{i_1 \cdots i_D}(t^w \cdot U)) = v(p_{i_1 \cdots i_D})(U) + \sum_{r=1}^D < \chi_i, w >$ is minimal. As a subspace is determined by its Plücker coordinates, we see that $(t^w \cdot U)_0$ is determined solely by the linear inequalities comparing the values $v(p_{i_1 \cdots i_D})(U) + \sum_{r=1}^D < \chi_i, w >$ as $(i_1, \ldots, i_D)$ varies. Thus, the equivalence classes form a polyhedral subdivision of $\mathbb{R}^n$ (and, in fact, a coherent one). In the constant coefficient case, all of the $v(p_{i_1 \ldots i_D})$ terms are 0 or ∞ so are the inequalities have no constant term and we get a fan.

$\square$

$\square$

The next lemma, which we will use repeatedly, says that the local geometry of $\operatorname{Trop} X$ can be reduced to the constant coefficient case.

Proposition 2.2.3. *Let $\kappa((t))^{\mathrm{alg}}$ denote the algebraic closure of the Laurent series field over κ. Let X be a subvariety of $(K^*)^n$ and $w \in \mathbb{R}^n$. Then the link of* $\operatorname{Trop} X$ *at w is* $\operatorname{Trop}\left(\operatorname{in}_w X \otimes_k \kappa((t))^{\mathrm{alg}}\right)$. *More specifically, for any $u \in \mathbb{R}^n$ we have* $\operatorname{in}_{w+\epsilon u} X = \operatorname{in}_u\left(\operatorname{in}_w X \otimes_k \kappa((t))^{\mathrm{alg}}\right)$ *for $\epsilon > 0$ and sufficiently small.*

Proof. The second claim implies the first because $u \in \operatorname{link}_w \operatorname{Trop} X$ if and only if $\operatorname{in}_{w+\epsilon u} X \neq \emptyset$ for $\epsilon > 0$ sufficiently small and $u \in \operatorname{Trop}\left(\operatorname{in}_w X \otimes_k \kappa((t))^{\mathrm{alg}}\right)$ if and only if $\operatorname{in}_u\left(\operatorname{in}_w X \otimes_k \kappa((t))^{\mathrm{alg}}\right) \neq \emptyset$.

As in the proof above, complete $(K^*)^n$ to projective space and abuse notation by writing the compactified X as $\operatorname{Proj} K[y_0, \ldots, y_n]/I$. Then the ideal $\operatorname{in}_u\left(\operatorname{in}_w X \otimes_k \kappa((t))^{\mathrm{alg}}\right)$ is finitely generated by elements of the form $\operatorname{in}_u \operatorname{in}_w f$ with $f \in I$. (Technically, this should be $\operatorname{in}_u \iota(\operatorname{in}_w f)$ where ι is the injection of $k[y_0, \ldots, y_N]$ into $\kappa((t))^{\mathrm{alg}}[y_0, \ldots, y_n]$.) Let $\operatorname{in}_u \operatorname{in}_w f_1, \ldots, \operatorname{in}_u \operatorname{in}_w f_m$ generate $\operatorname{in}_u \operatorname{in}_w I$, for $f_j \in I$. Then, for ϵ small enough, we will have $\operatorname{in}_{w+\epsilon u} f_j = \operatorname{in}_u \operatorname{in}_w f_j$ for every j. Thus, for such an ϵ we have $\operatorname{in}_{w+\epsilon u} I \supseteq \operatorname{in}_u \operatorname{in}_w I$. But both ideals have the same Hilbert function, so they are equal. $\square$

Proposition 2.2.4. *Suppose $w \in \operatorname{Trop} X$ is contained in the relative interior of a face σ of Σ, where Σ is adapted to $\operatorname{Trop} X$. Let $H(\sigma)$ be the translation to the origin of the affine linear space spanned by σ. Then $\operatorname{in}_w X$ is invariant under translation by the torus $\exp(H(\sigma))$.*

Proof. Let $v \in H(\sigma)$, then for small ϵ we have $\operatorname{in}_{w+\epsilon v} X = \operatorname{in}_w X$. Using the previous lemma, we see that, for any $f \in \operatorname{in}_w I$, we also have $\operatorname{in}_v f = \operatorname{in}_v \operatorname{in}_w f = \operatorname{in}_{w+\epsilon v} f = \operatorname{in}_w f$. Thus, the terms of f which have lowest weight in the v-grading again lie in $\operatorname{in}_w I$. Subtracting them off and repeating, we may conclude that $\operatorname{in}_w I$ is homogeneous with respect to the v grading. This is equivalent to being invariant under the torus $\exp(\mathbb{R}v)$. $\square$

2.3 Degenerating Toric Varieties

The construction in this section appears to have been discovered several times, the earliest reference I can find is [35]. We repeat it here as it does not appear to be well known. Let Σ be a (finite) polyhedral complex in $\mathbb{R}^n$ whose faces are defined by inequalities of the form $\sum a_i x_i \geq w$ with $a_i \in \mathbb{Z}$ and $w \in v(K^*)$. Define Σ_1 to be the fan whose cones are the recession cones of the faces of Σ. We will construct, associated to Σ, a flat family $\mathcal{T}$ over $\operatorname{Spec} \mathcal{R}$ whose fiber over $\operatorname{Spec} K$ is the toric variety $\overline{T}$ assosciated to Σ_1 and whose fiber over $\operatorname{Spec} \kappa$ is a union of toric varieties $\overline{T}_0$ indexed by the faces of Σ. As a running example, we will take the case where Σ is the complex in $\mathbb{R}^2$ shown in Figure 2.1. $\overline{T}$ will be

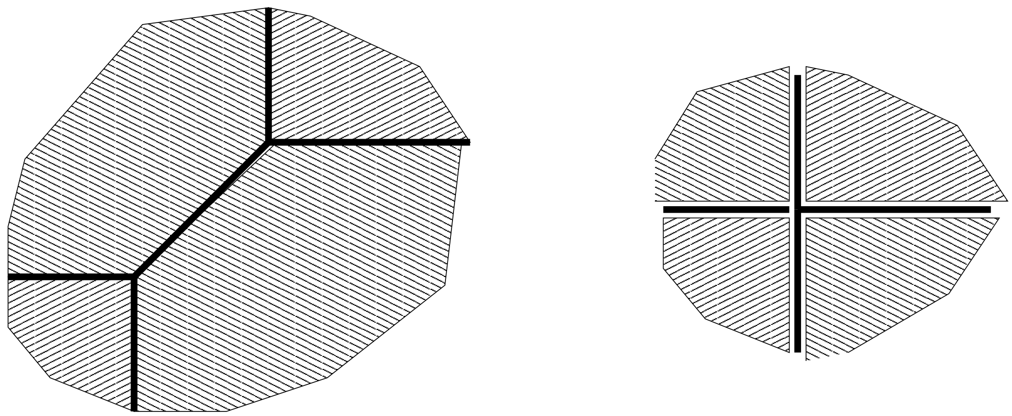

Figure 2.1: Σ and Σ_1 for our running example

$\mathbb{P}^1 \times \mathbb{P}^1$ and $\overline{T}_0$ will be the union of two $\mathbb{P}^2$'s glued along a $\mathbb{P}^1$.

This construction is better known in the case where Σ is dual to a regular subdivision of a polytope. In that case, this construction appears in [43] and is used in Viro's patchworking construction [47]. Our running example is of this sort, with Σ dual to the subdivision of the square $\mathrm{Hull}((0,0),(1,0),(1,1),(0,1))$ into the triangles $\mathrm{Hull}((0,0),(1,0),(1,1))$ and $\mathrm{Hull}((1,0),(1,1),(0,1))$. When Σ is dual to a subdivided polytope in this manner, the degeneration of $\overline{T}$ to $\overline{T}_0$ takes place within a projective space. In our example, the family $\overline{\mathcal{T}}$ can be embedded in $\mathbb{P}^3_{\mathcal{R}}$ as $wz = txy$.

Let σ be a face of Σ. Define $R(\sigma)$ to be the following subring of $K[x_1^{\pm}, \ldots, x_n^{\pm}]$: the elements of $R(\sigma)$ are those which can be expressed as sums $\sum a_I x^I$ such that, for every I and for every $w \in \sigma$, we have $< w, I > +v(a_i) \geq 0$. Whenever τ is a face of σ, we have a natural inclusion $R(\sigma) \hookrightarrow R(\tau)$ and thus a natural map $\mathrm{Spec}\, R(\tau) \to \mathrm{Spec}\, R(\sigma)$ which turns out to be an inclusion. Gluing the latter as in the standard construction of a toric variety, we build $\mathcal{T}$.

In our example, the vertices at $(0,0)$ and $(1,1)$ correspond to the rings $\mathcal{R}[x_1^{\pm}, x_2^{\pm}]$ and $\mathcal{R}[(t^{-1}x_1)^{\pm}, (t^{-1}x_2)^{\pm}]$ respectively. The spectrum of each of these rings is a flat family over $\mathrm{Spec}\, \mathcal{R}$ with general and special fibers each n-dimensional tori. Because $(0,0)$ and $(1,1)$ both correspond to the same face of Σ_1 (the origin) their general fibers are glued together, but their special fibers, which correspond to different faces of Σ, are not. The edge running from $(0,0)$ to $(1,1)$, which we will denote e, corresponds to the ring $\mathcal{R}[x_1^{\pm}, x_2^{\pm}] \cap \mathcal{R}[(t^{-1}x_1)^{\pm}, (t^{-1}x_2)^{\pm}] = \mathcal{R}[(x_1 x_2^{-1})^{\pm}, x_1, t x_1^{-1}]$. This can also be written as $\mathcal{R}[u^{\pm}, v, w]/(vw - t)$. This ring corresponds to a family over $\mathrm{Spec}\, \mathcal{R}$ whose fiber over K is $\mathrm{Spec}\, K[u^{\pm}, v^{\pm}] = (K^*)^n$ and whose fiber over κ is $\mathrm{Spec}\, \kappa[u^{\pm}, v, w]/(vw)$, which is two copies of $\kappa^* \times \kappa$ glued along a κ^*. The inclusions of the two endpoints of this segment into this segment correspond to inclusions of $R((0,0))$ and $R((1,1))$ into $R(e)$. In each of these inclusions, the map on fibers over K is an isomorphism and the map on fibers over κ takes $(\kappa^*)^2$ into one of the two copies of $\kappa^* \times \kappa$. Adding in the patches from the other faces of Σ, we get a family whose fiber over K is $\mathbb{P}^1_K \times \mathbb{P}^1_K$ and whose fiber over κ is two copies of $\mathbb{P}^2_\kappa$ glued along a $\mathbb{P}^1_\kappa$.

We now return to discussing our general construction. Let $\overline{T}$ denote $\mathcal{T} \otimes_{\mathcal{R}} K$ and $\overline{T}_0$ denote $\mathcal{T} \otimes_{\mathcal{R}} \kappa$. Using $\partial \overline{T}$ for the toric boundary of $\overline{T}$ and $\partial \overline{T}_0$ for its closure in $\overline{T}_0$, let $T_0 = \overline{T}_0 \setminus \partial T_0$. So T_0 is a flat degeneration of the torus $(K^*)^n$ and $\overline{T}$ and $\overline{T}_0$ are partial compactifications of $(K^*)^n$ and T_0 respectively. We will sometimes denote $(K^*)^n$ as T to be consistent with the rest of this notation.

For $\sigma \subset \mathbb{R}^n$ a polytope, let $H(\sigma)$ denote the translation to the origin of the affine linear space spanned by σ. There are natural actions of $(K^*)^n$ and $(\kappa^*)^n$ on $\overline{T}$ and $\overline{T}_0$ respectively such that the orbits are in bijection with the faces of Σ_1 and Σ respectively. This correspondence is inclusion reversing on the closures and the orbit $\mathcal{O}_\sigma$ corresponding to a face σ has stabilizer $\exp H(\sigma)$. Note that, as with a standard toric variety, the following potentially confusing point exists: each face $\sigma \in \Sigma$ corresponds to a coordinate patch $\operatorname{Spec} R(\sigma) \otimes_{\mathcal{R}} \kappa$ on $\overline{T_0}$. This patch is $\bigcup_{\tau \subseteq \sigma} \mathcal{O}_\tau$. The assignment of patches to faces is inclusion preserving; that of orbits to faces is inclusion reversing.

We will need the following lemma often in the future:

Lemma 2.3.1. *Let $w \in v(K^*)^n$ and suppose that w lies in the relative interior of $\sigma \subset \mathbb{R}^n$ where σ is a face of Σ. Consider the point $t^w \in (K^*)^n \subset \mathcal{T}$. Then the limit of t^w in $\overline{T}_0$ lies in the torus orbit corresponding to σ.*

Proof. We may do this computation while looking solely at the coordinate patch $\operatorname{Spec} R(\sigma)$. Define a map $R(\sigma) \to \mathcal{R}$ by $\sum a_I x^I \mapsto \sum a_I t^{<w,I>}$. By the definition of $R(\sigma)$ and the assumption that $w \in \sigma$, the image truly lands in $\mathcal{R}$ (and not just K). By definition, over K, this is the inclusion of the point t^W into $(K^*)^n$. Thus, the limit of t^w does lie somewhere in $\operatorname{Spec} R(\sigma) \otimes_{\mathcal{R}} \kappa$ and must lie in the orbit $\mathcal{O}_\tau$ for some $\tau \subseteq \sigma$. Thus it is enough to show that the limit of t^w is **not** in the coordinate patch corresponding to τ for any $\tau \supsetneq \sigma$.

Let $\tau \supsetneq \sigma$. Since w is in the relative interior of σ, we know that $w \notin \tau$ and we can find an affine linear functional $\sum a_i x_i + v$ which is 0 on τ but negative at w. Then $t^v x_1^{a_1} \cdots x_n^{a_n}$ is a monomial that is in $R(\tau)$. Evaluating this monomial at $x_i = t^{w_i}$ produces a negative power of t and thus not a member of $\mathcal{R}$. So the map $R(\sigma) \to \mathcal{R}$ by evaluation at $x_i = t^{w_i}$ can not be extended to a map $R(\tau) \to \mathcal{R}$. The corresponding geometric statement is that the map $\operatorname{Spec} \mathcal{R} \hookrightarrow \operatorname{Spec} R(\sigma)$ which sends $\operatorname{Spec} K$ to t^w does not factor through $\operatorname{Spec} R(\tau)$.

□

2.4 The Tropical Degeneration and Compactifications

Let $X \subset (K^*)^n$ and let $\Sigma \subset \mathbb{R}^n$ be a polyhedral complex. In the previous section, we defined the family $\mathcal{T}$ over $\operatorname{Spec} \mathcal{R}$, with generic fiber a toric variety $\overline{T}$ and special fiber $\overline{T}_0$. We define $\overline{X}$ to be the closure of X in $\overline{T}$, $\overline{X}_0$ to be the closure of X in $\overline{T}_0$ and X_0 be the closure of X in T_0; we have $X_0 = \overline{X}_0 \cap T_0$. We will refer to $\overline{X}$, X_0 and $\overline{X_0}$ as the "tropical compactification", "tropical degeneration" and "compactified tropical degeneration" of X, respectively. These constructions make sense for any Σ, but their importance arise when Σ is supported on $\operatorname{Trop} X$ and sufficiently fine. In this case, we will see that $\overline{X}_0$ is covered by strata indexed by the faces of Σ and isomorphic to quotients by tori of the various $\operatorname{in}_w X$.

Before continuing with the general theory, we pause to work out the example of the hypersurface X given by $t+x+y+xy=0$ in $(K^*)^2$. The special case where X is a hypersurface is central in the Patchworking construction of Viro ([47]) and in the work of Gelfand, Kapranov and Zelevinsky ([21]). Sturmfels has generalized Viro's work to complete intersections in [44].

With $X = \{(x,y) : t+x+y+xy=0\}$, it is easy to check that $\operatorname{Trop} X$ is the one skeleton of the polyhedral complex Σ in Figure 2.1. We saw in the previous section that the polyhedral complex in Figure 2.1 corresponded to a copy of $\mathbb{P}^1 \times \mathbb{P}^1$ over K degenerating to two copies of $\mathbb{P}^2$ over κ. Taking the 1-skeleton of Σ corresponds to removing the four torus fixed points from each fiber. Let e denote the diagonal edge of Σ. The edge e corresponds, as we saw in the last section, to a coordinate patch $\operatorname{Spec} \mathcal{R}[(xy^{-1})^{\pm}, x, tx^{-1}] = \operatorname{Spec} \mathcal{R}[u^{\pm}, v, w]/(vw - t)$. The ideal of $\mathcal{X}$ inside this coordinate patch is found by intersecting the ideal $(t+x+y+zx)$ in $K[x^{\pm}, y^{\pm}]$ with $\mathcal{R}[(xy^{-1})^{\pm}, x, tx^{-1}]$. The resulting ideal is generated by $x^{-1}(t+x+y+xy) = tx^{-1} + 1 + (xy^{-1})^{-1} + x(xy^{-1})^{-1}$. In terms of the (u,v,w) variables, this is $v + 1 + u^{-1} + wu^{-1}$.

Consider the coordinate patch $\operatorname{Spec} \mathcal{R}[u^{\pm}, v, w]/(vw-t)$; its fiber over κ is $\operatorname{Spec} \kappa[u^{\pm}, v, w]/(vw)$. Geometrically, this is two copies of $\kappa \times \kappa^*$ glued along a κ^*. The ideal of $\overline{X}_0$ inside this coordinate patch (actually, only X_0 meets this patch) is cut out by $v+1+u^{-1}+wu^{-1}$. This cuts out a rational curve in each component of $\operatorname{Spec} \kappa[u^{\pm}, v, w]/(vw)$, given by $v+1+u^{-1}=0$ and $1+u^{-1}+wu^{-1}$

respectively. Note that the intersections of these curves with $(\kappa^*)^2$ are $\mathrm{in}_{(0,0)}$ and $\mathrm{in}_{(1,1)} X$. The intersection of $\overline{X}_0$ with the orbit $\mathcal{O}_e = \operatorname{Spec} \kappa[u^{\pm}]$ is given by the ideal $(u^{-1}+1)$. Note that this is the quotient of $\mathrm{in}_{(s,s)} X$ by its invariant torus for any $s \in e$.

This construction was discovered by Tevelev in the constant coefficient case, see [45], and also partially and in a messier form by myself. Hacking realized that this construction could be used to study the cohomology of $\mathrm{Trop}(X) \setminus \{0\}$. It seems not to have been defined before for the nonconstant coefficient case. This section has benefited greatly from conversations with P. Hacking.

Proposition 2.4.1. *X_0 and $\overline{X}_0$ are flat degenerations of X and $\overline{X}$ respectively. If we assume that the support of Σ contains $\operatorname{Trop} X$, then $\overline{X}$ and $\overline{X_0}$ are proper over K and κ respectively.*

Proof. X_0 and $\overline{X_0}$ are both defined as the fibers over $\operatorname{Spec} \kappa$ of the closures of X and $\overline{X}$ within the flat families $\mathcal{T} \setminus \partial\mathcal{T}$ and $\mathcal{T}$; this proves the first claim. Complete Σ to a polyhedral complex Σ' whose support is all of $\mathbb{R}^n$; let $\overline{T}'$ *etc.* denote the associated objects. $\overline{T}'$ is a proper toric variety, as it is associated to a complete fan and, similarly, $\overline{T}'_0$ is proper because it is a union of toric varieties each of which is given by a complete fan. $\overline{X}'$ and $\overline{X}'_0$ are similarly proper because they are closed subvarieties of $\overline{T}'$ and $\overline{T}'_0$ respectively.

We will show that, in fact, $\overline{X}' = \overline{X}$ and $\overline{X}'_0 = \overline{X}_0$. Let σ be a face of Σ' not in Σ and let $\mathcal{O}_\sigma$ be the corresponding orbit in $\overline{T}'_0$. We will show that the closure of X is disjoint from $\mathcal{O}_\sigma$. Repeating this for every such σ shows that $\overline{X}'_0$ in fact lies entirely in $\overline{T}_0$ and thus is $\overline{X}_0$.

Every point of the closure of X can be approached along a one dimensional path through X. More precisely, let D be $\operatorname{Spec} S$ for S some discrete valuation ring with fraction field L, residue field λ and uniformizer u. Let $\phi : D \to \overline{\mathcal{T}}'$ with $\phi(\operatorname{Spec} L) \in X$. We want to show that $\phi(\lambda) \notin \mathcal{O}_\sigma$; we will then know that the closure of X is disjoint from $\mathcal{O}_\sigma$. If $\phi(\lambda)$ lands in the fiber above $\operatorname{Spec} K$ then trivially it is not in $\mathcal{O}_\sigma$, so we are only interested in the case where $\phi(\lambda)$ lands above κ. In this case, the projection from $\overline{\mathcal{T}}' \to \operatorname{Spec} \mathcal{R}$ gives us a surjective map $\operatorname{Spec} S \to \operatorname{Spec} \mathcal{R}$ and we can extend v to a map $L^* \to \mathbb{R}$, which we will also denote as v. Writing x_i for the coordinate functions on the big torus in which X lives, we may now consider $w := (v(\phi^*(x_1)), \ldots, v(\phi^*(x_n)))$.We have $w \in \operatorname{Trop} X$, as definition 4 of $\operatorname{Trop} X$ in Theorem 2.1.2 is clearly invariant under extending the ground field to L. Then w is not in the relative interior of σ, as $\sigma \notin \Sigma$ and $\operatorname{Trop} X$ is supported on Σ. So, by Lemma 2.3.1, $\phi(\lambda)$ is not in $\mathcal{O}_\sigma$. We now know that $\overline{X}'_0 = \overline{X}_0$ as promised and is proper.

One could use a similar argument to see that $\overline{X}' = \overline{X}$ and is thus proper. A simpler argument is to see that $D := \overline{\mathcal{X}}' \setminus \overline{\mathcal{X}}$ is closed. (It is the intersection of $\overline{\mathcal{X}}$ with the closed subvariety of $\mathcal{T}$ corresponding to the faces of Σ' and Σ'_1 not it in σ and Σ'.) Since $\overline{\mathcal{X}}'$ is proper over $\operatorname{Spec} \mathcal{R}$ the image of D in $\operatorname{Spec} \mathcal{R}$ must be closed. But we have just seen that $\operatorname{Spec} \kappa$ is not the image of D, so the image of

D is empty and thus $D = \emptyset$ and $\overline{X} = \overline{X}'$.

$\square$

The idea of this paragraph is due to Tevelev and Hacking in the constant coefficient case: Construct a family $\mathcal{F} \subset \overline{\mathcal{T}} \times_{\mathcal{R}} \operatorname{Spec} \mathcal{R}[x_1^{\pm}, \ldots, x_n^{\pm}]$ over $\overline{\mathcal{T}}$ as follows: Over $x \in (K^*)^n$ the fiber $\mathcal{F}_x$ is $x^{-1} \cdot X$. $\mathcal{F}$ is then the closure of this family in the n-torus family over $\overline{\mathcal{T}}$. In this setting, $\mathcal{X}$ can be described as $\{x \in \overline{\mathcal{T}} : (1, \ldots, 1) \in \mathcal{F}_x\}$. We now show

Theorem 2.4.2. *There is a choice of polyhedral structure Σ on $\operatorname{Trop} X$ such that $\mathcal{F}$ is flat over $\mathcal{X}$; for the rest of the statement of this theorem assume that Σ has this property. Let $x \in \mathcal{O}_\sigma$ for $\sigma \in \Sigma$. $\mathcal{O}_\sigma$ contains a canonical point x_0 which is the flat limit of t^w for every w in the relative interior of σ. Let $x = sx_0$ for $s \in (\kappa^*)^n$, then the fiber of $\mathcal{F}$ over x is $s^{-1} \cdot \operatorname{in}_w X$.*

Proof. Let $I \subset K[x^{\pm}]$ be the ideal of X and take a provisional choice of Σ fine enough to be adapted to X.

Lemma 2.4.3. *Fix $\sigma \in \Sigma$ and let $J = \operatorname{in}_w I$ for w in the relative interior of σ. Let f be a polynomial in J which is homogeneous and of degree 0 with respect to the action of $\exp(H(\sigma))$. There is a finite collection $g_1, \ldots, g_r$ of members of I such that, for every w in the interior of σ we have $\operatorname{in}_w g_i = f$ for some $1 \leq i \leq r$.*

In the constant coefficient case, we could simply take Σ fine enough to be adapted to the closure of X in $\mathbb{P}^n$. Then we could construct such g_i from a universal Gröbner basis. We prefer to give a direct proof rather than to adapt the proof of the existance of a universal Gröbner basis to the non-constant coefficient case.

Proof of Lemma 2.4.3. We first summarize the strategy of our proof. Consider some w in the relative interior of σ. There is a $g_w \in I$ such that $\mathrm{in}_w\, g_w = f$. Moreover, there will be an open subset U_w of σ containing w such that $\mathrm{in}_{w'}\, g_w = f$ for $w' \in U$. If the relative interior of σ were compact, we would then be able to take a finite number of U_w's covering the relative interior of σ and the assosciated g_w's would satisfy the claim.

Unfortunately, the relative interior of σ is not compact. Therefore, we impose a condition on w' which is a bit more complicated than simply asking that $\mathrm{in}_{w'}\, g_w = f$ so that we can work with w on the bondary of σ. Secondly, σ may also not be compact. We counter this by taking the cone on σ and intersecting it with the unit sphere. It is simplest to present both modifications at once.

Embed σ into $\mathbb{R}^n \times \mathbb{R}_{\geq 0}$ and let $\tilde{\sigma}$ be the closed cone over σ. Let $(w, e) \in \mathbb{R}^n \times \mathbb{R}_{\geq 0}$ and let $f = \sum_{E \in A} f_E X^E \in K[x^{\pm}]$. We define $\mathrm{in}_{(w,e)}\, f$ as follows: for $e > 0$ we put $\mathrm{in}_{(w,e)}\, f = \mathrm{in}_{(w/e)}\, f$ and for $e = 0$ we put $\mathrm{in}_{(w,0)}(\sum_{I \in A} f_E x^E) = \sum_{E \in B} f_E x^E$ where B is the subset of A on which $< w, E >$ is maximized. Note that $\mathrm{in}_{(w,e)}\, f \in \kappa[x^{\pm}]$ for $e > 0$ and $\mathrm{in}_{(w,0)}\, f \in K[x^{\pm}]$. Note that, given $f \in K[x^{\pm}]$ and $(w_1, e_1), (w_2, e_2) \in \mathbb{R}^n \times \mathbb{R}_{\geq 0}$, we have $\mathrm{in}_{(w_1,e_1)}\, \mathrm{in}_{(w_2,e_2)}\, f = \mathrm{in}_{(\epsilon w_1 + w_2, \epsilon e_1 + e_2)}\, f$ for $\epsilon > 0$ sufficiently small. (Here, when $e_2 > 0$, the left hand side is techinically

$\mathrm{in}_{(w_1,e_1)} \iota(\mathrm{in}_{(w_2,e_2)} f)$ where ι is the obvious embedding $\kappa[x^{\pm}] \hookrightarrow \kappa((t))^{\mathrm{alg}}[x^{\pm}]$.) Also, note that we have $\mathrm{in}_{(w,e)} f = \mathrm{in}_{(\lambda w, \lambda e)} f$ for $\lambda \in \mathbb{R}_{>0}$.

Our proof is by induction on $D = \dim \sigma + \epsilon$ where $\epsilon = 0$ if we are in the constant coefficients case and 1 otherwise.

Let $(w,e) \in \tilde{\sigma}$, $(w,e) \neq (0,0)$. We claim that there is an open neighborhood U of (w,e) in $\tilde{\sigma}$ a finite number of polynomials $g_1, \ldots, g_r$ such that, for every (w',e') in the intersection of U and the relative interior of $\tilde{\sigma}$, we have $\mathrm{in}_{(w',e')} g_i = f$ for some g_i. Let us see why this finishes the proof: Since $\mathrm{in}_{(w',e')}$ is invariant under scalaing (w',e'), we may assume that U is homothety invariant. The intersection of $\tilde{\sigma}$ with the sphere of radius 1 is compact, so we may cover $\tilde{\sigma}$ with a finite number of these U. Then the union of the finite number of finite collections of g's meets the conditions of the theorem.

We now must prove the claim of the previous paragraph. Let τ be the face of $\tilde{\sigma}$ containing (w,e). We have $\dim \tau \geq 1$. Let $I_\tau = \mathrm{in}_{(w,e)} I$. Then I_τ is homogenous with respect to $\exp(H(\tau) \cap \mathbb{R}^n)$. Let Q denote the cone of vectors (u,d) such that $(w,e) + \epsilon(u,d) \in \tilde{\sigma}$ for $\epsilon > 0$ small enough. The ideal I_τ is defined over the ring R of polynomials homogenous with respect to $\exp(H(\tau) \cap \mathbb{R}^n)$; write I' for $I_\tau \cap R$. (R might be either a polynomial ring defined over κ or over K depending on whether τ is in $\mathbb{R}^n \times \{0\}$ or not, we are trying to emphasize the similarity of the two cases.) Since J is homogenous with respect to $\exp(H(\sigma)) \supseteq \exp(H(\tau) \cap \mathbb{R}^n)$, J' is also generated over R; let J' be $J \cap R$.

We want to apply our inductive hypothesis with I', J' and $Q/(H(\tau)\cap\mathbb{R}^n)$ in place of I, J and σ. If τ is contained in $\mathbb{R}^n$ then $H(\tau)\cap\mathbb{R}^n$ is at least 1 dimensional so transferring our attention to $Q/(H(\tau)\cap\mathbb{R}^n)$ reduces D by 1. If we are not already in the constant coefficients case when considering I then we are when considering I' so again D is reduced by 1. Finally, if we were already in the constant coefficients case when considering I then Σ was a fan and everything is invariant under scaling $\mathbb{R}^n\times\mathbb{R}_{\geq 0}$ both only along the first n coordinates and only along the last coordinate. This allows us to replace the use of the unit sphere by the unit sphere in $\mathbb{R}^n\times\{1\}$, which only meets τ when $\tau\cap\mathbb{R}^n$ has dimension at least 1 (*i.e.* τ is not the cone on the vertex of σ), so again D goes down by one.

We conclude that there is a finite number of polynomials $g_1^\tau, \ldots, g_r^\tau \in I'$ such that, for every (u,d) in the relative interior of Q, $\mathrm{in}_{(u,d)}\, g_i^\tau = f$ for some g_i^τ. Since $g_i^\tau \in I' \subset I^\tau$, each g_i is $\mathrm{in}_{(w,e)}\, g_i$ for some $g_i \in I$.

We first see that there is an open set $U^\tau \subset \tau$ containing (w,e) such that $\mathrm{in}_{(w',e')}\, g_i = g_i^\tau$ for $(w',e') \in U^\tau$. This is simple enough: writing $g_i = \sum_{E\in A} g_{E,i}x^A$, g_i^τ is a sum over the subset of A on which $ev(g_{E,i})+ < w, E >$ is minimized, call this subset B. We know that g_I^τ is homogenous for $\exp(H(\tau)\cap\mathbb{R}^n)$, which means that $< w', E >$ is constant on B for $w' \in H(\tau)\cap\mathbb{R}^n$. Thus, as (w',e') moves through τ, $e'v(g_{E,i})+ < w', E >$ remains constant on B. We just have to make sure that this constant value is still the minimum, which amounts to a finite list of inequalities.

By shrinking U^τ, we may arrange that the closure of U^τ is compact, contained in the relative interior of τ and that $\mathrm{in}_{w',e'} g_i = g_i^\tau$ also holds for (w', e') in this closure. Now, for every (u, d) in the relative interior of τ and every $(w', e') \in U^\tau$, there is an i such that $\mathrm{in}_{(u,d)} \mathrm{in}_{(w',e')} g_i = f$. For $\epsilon > 0$ small enough , we have

$$\mathrm{in}_{(u,d)} \mathrm{in}_{(w',e')} g_i = \mathrm{in}_{\epsilon(u,d)+(w',e')} g_i.$$

Let us restrict (u, d) to range over those vectors in the relative interior of τ whose components are all less than 1. We claim that we can choose ϵ uniformly, independent of (u, d) (subject to the preceeding restriction) and (w', e') so that the displayed equality holds.

Let $g_i = \sum_{E \in A} g_{E,i} x^E$. Let B be the subset of A on which $e'v(g_{E,i})+ < w', E >$ is minimized for $(w', e') \in U^\tau$. Let C be the further subset of B on which $dv(g_{E,i})+ < u, E >$ is minimized. Then $\mathrm{in}_{(u,d)} \mathrm{in}_{(w',e')} g_i$ is a sum over C. We will have $\mathrm{in}_{\epsilon(u,d)+(w',e')} g_i = \mathrm{in}_{(u,d)} \mathrm{in}_{(w',e')} g_i$ as long as C is also the subset of A on which $(\epsilon d + e')g_i$ as long as C is also the subset of A on which

$$(\epsilon d+e')v(g_{E,i})+ < \epsilon u+e', E >= \epsilon(dv(g_{E,i})+ < u, E >)+(e'v(g_{E,i})+ < w', E >)$$

assumes its minimum; call this subset C'. Now, $e'v(g_{E,i})+ < w', E >$ is consant on B, so $C' \cap B = C$. Thus, we simply want to insure that $C' \subseteq B$ or, in other words, that $|\epsilon(dv(g_{E,i})+ < u, E >)|$ is always less than the difference between $e'v(g_{E,i})+ < w', E >$ on B and $e'v(g_{E,i})+ < w', E >$ on $A \setminus B$. Now, we have required that $\mathrm{in}_{(w',e')} g_i = g_i^\tau$ for (w', e') in the closure of U^τ, so the value of

$e'v(g_{E,i})+ < w', E >$ on B is greater than its value on $A \setminus B$ for all (w', e') in this closure. Since we took the closure of U^τ to be compact, there is some positive lower bound for the difference between $e'v(g_{E,i})+ < w', E >$ on B and $e'v(g_{E,i})+ < w', E >$ on $A \setminus B$ for all $(w', e') \in U^\tau$. On the other hand, since we took (u, d) to be bounded, there is an upper bound for $dv(g_{E,i})+ < u, E >$ and we see that we can indeed choose ϵ uniformly.

Let V denote the subset of the relative interior of τ where all coordinates are less than 1. $U^\tau + \epsilon V$ is the desired U. $\square$

We now continue with our proof of Theorem 2.4.2. For every face $\sigma \in \Sigma$ and w in the relative interior of σ take a finite generating set $f_1, \ldots, f_r$ for $\mathrm{in}_w I$. By the lemma, for each f_i we can find a finite set of polynomials $g_i^1, \ldots, g_i^{s_i}$ so that, for every w in the relative interior of σ we have $\mathrm{in}_w g_i^t = f_i$ for some $1 \leq t \leq s_i$. By refining Σ, we may guarantee that the following holds: for every $\sigma \in \Sigma$ there is a generating set f_i^σ of $\mathrm{in}_w I$ (for w in the relative interior of σ) and a collection of polynomials g_i^σ such that $\mathrm{in}_w g_i^\sigma = f_i^\sigma$ for all w in the relative interior of σ. Moreover, we can require that f_i^σ be homogeneous with respect to $\exp(H(\sigma))$, that the coefficient of x^0 in f_i^σ be nonzero (by multiplying by a polynomial) and that the valuation of the coefficient of x^0 in g_i^σ be 0 (by multiplying by a scalar).

Let $g_i^\sigma = \sum_{E \in A} g_E^{\sigma,i} x^E$ and $f_i^\sigma = \sum_{E \in A'} f_E^{\sigma,i} x^E$. By the normalizations at the end of the above paragraph, for every w in the relative interior of σ we have

$v(g_E^{\sigma,i}) + < w, E > \geq 0$ and achieves the minimal value of 0. This implies that $g_i^\sigma \in R(\sigma)$. Now, let $x \in \mathcal{O}_\sigma$ and let x_0 and s be as in the statement of the theorem. We will now prove that the fiber of $\mathcal{F}$ over x is $s^{-1} \cdot \mathrm{in}_w X$. We may compute in the coordinate patch $\operatorname{Spec} R(\sigma)[y_1^\pm, \ldots, y_n^\pm]$. The polynomial $\sum g_E^{\sigma,i} x^E y^E$ is in the ideal defining $\mathcal{F}$ inside $(K^*)^n$ and lies in $\operatorname{Spec} R(\sigma)[y_1^\pm, \ldots, y_n^\pm]$ by the preceding, so it vanishes on the restriction of $\mathcal{F}$ to $\operatorname{Spec} R(\sigma)[y_1^\pm, \ldots, y_n^\pm]$. Specializing to $y_i = s_i$ and setting $\mathcal{M} = 0$, we get that $\sum f_E^{\sigma,i} s^E x^E$ vanishes on the fiber of $\mathcal{F}$ over x. Since the polynomials $\sum f_E^{\sigma,i} s^E x^E$ generate $s^{-1} \cdot \mathrm{in}_w I$, we see that $\mathcal{F}_x \subseteq s^{-1} \mathrm{in}_w I$. On the other hand, $s^{-1} \cdot \mathrm{in}_w X$ is the flat limit of the restriction of $\mathcal{F}$ to the one parameter family st^w, so $\mathcal{F}_x$ must contain this limit and we see that $\mathcal{F}_x = s^{-1} \cdot \mathrm{in}_w X$.

Finally, we must prove flatness of $\mathcal{F}$. As flatness is an open condition on the base, it is enough to show that $\mathcal{F}$ is flat over $\overline{X_0}$. Take $x \in \overline{T_0}$. By Theorem 4.2.8.a of [32] it is enough to find a collection of maps $\operatorname{Spec} S_i \to \overline{\mathcal{T}}$ from the spectra of discrete valuation rings to $\overline{\mathcal{T}}$ with the closed point of each $\operatorname{Spec} S_i$ mapping to x such that (1) the map $A \hookrightarrow \prod S_i$ from the local coordinate ring of $\overline{\mathcal{T}}$ at x to the product of the S_i is injective and (2) $\mathcal{F} \times_{\overline{\mathcal{T}}} \operatorname{Spec} S_i$ is flat over $\operatorname{Spec} S_i$. The collection of sections $\operatorname{Spec} \mathcal{R} \to \overline{\mathcal{T}}$ with $\operatorname{Spec} K$ landing in $(K^*)^n$ and $\operatorname{Spec} \kappa$ hitting x obeys the first condition.

The map $\operatorname{Spec} K \to \overline{T}$ is given by an n-tuple $(u_1, \ldots, u_n) \in (K^*)^n$. Let $w_i = v(u_i)$. The assumption that $\operatorname{Spec} \kappa$ is taken to a point of $\mathcal{O}_\sigma$ implies that

$(w_1, \ldots, w_n) := (v(u_1), \ldots, v(u_n))$ is in the relative interior of σ and that, writing $u_i = \ell_i t^{w_i} \mod t^{w_i}\mathcal{M}$ for $\ell_i \in \kappa$, we have $(\ell_1, \ldots, \ell_n) \cdot x_0 = x$. The flat limit of $\mathcal{F}$ along this family is $(\ell_1, \ldots, \ell_n)^{-1}$ times the limit along t^w, which is $\mathrm{in}_w X$. This completes the proof of Theorem 2.4.2. □

Set $X_\sigma = \overline{X_0} \cap \mathcal{O}_\sigma$, so $\overline{X_0} = \bigsqcup X_\sigma$. We now list the main properties of the X_σ.

Proposition 2.4.4. *1. If X_σ is in the closure of X_τ then $\tau \subset \sigma$.*

2. For any σ, the union $\bigcup_{\tau \subseteq \sigma} X_\sigma$ is affine.

3. For any σ, the union $\bigcup_{\tau \supseteq \sigma} X_\sigma$ is proper.

4. Assume X is irreducible. Then X_σ is $d - \dim \sigma$ dimensional, where d is the dimension of X.

Proof. (1) This just comes from the closure relation on the orbits $\mathcal{O}_\sigma$.

(2) We first note that $\bigcup_{\tau \subseteq \sigma} \mathcal{O}_\sigma$ is affine by construction. Then $\bigcup_{\tau \subseteq \sigma} X_\sigma$ is a closed subvariety of that affine variety.

(3) We use the same trick as in the proof of the previous proposition. First, complete Σ to a complete polyhedral complex Σ'. Then $\bigcup_{\tau \supseteq \sigma'} \mathcal{O}_{\sigma'}$ is proper when the union is taken over $\sigma' \in \Sigma'$ and $\bigcup_{\tau \supseteq \sigma} X'_\sigma$ is a closed subvariety of this proper variety and hence is proper. But, if σ is a face of Σ' not in Σ, then X'_σ is empty.

(4) Let w lie in the relative interior of σ, so by Proposition 2.2.4, $\mathrm{in}_w X$ is invariant under $\exp(H(\sigma))$. We claim that $X_\sigma = \mathrm{in}_w X/\exp(H(\sigma))$. We show that $\mathrm{in}_w X$ has dimension d, so the dimension of $X_\sigma = d - \dim H(\sigma)$. Proof that $\mathrm{in}_w X$ has dimension d: $\mathrm{in}_w X$ is the intersection in $\mathrm{Spec}\,\mathcal{R}[x^{\pm}]$ of the closure of $t^w X$ and the fiber over $\mathrm{Spec}\,\kappa$. The closure of $t^w X$, henceforth denoted $\overline{t^w X}$, is of dimension $d+1$ and is irreducible since it is the closure of an irreducible scheme $t^w X$. So the intersection of $\overline{t^w X}$ with the fiber over $\mathrm{Spec}\, m$ has dimension at least d, and we must have equality because $\overline{t^w X}$ also lies over the general fiber. (This argument is a modification of [22], first pararaph of the proof of Theorem 1.) We now check the claim.

Every point $x \in \mathcal{O}_\sigma$ is of the form $x = s \cdot x_0$ where x_0 is the limit of t^w and $s \in (\kappa^*)^n$. Note that s is determined only modulo $\exp(H(\sigma))$. In the notation of the preceding lemmas, we have $\mathcal{F}_x = s^{-1} \cdot \mathrm{in}_w X$ and $x \in X_\sigma$ if and only if $(1, \ldots, 1) \in \mathcal{F}_x$. We thus see that $x \in X_\sigma$ if and only if $s \in \mathrm{in}_w X$, so $X_\sigma \cong \mathrm{in}_w X/\exp(H(\sigma))$. □

Remark: We will not use the family $\mathcal{F}$ again, but it is very powerful. Hacking, has observed that, if every $\mathrm{in}_w X$ is smooth as a scheme, then it follows that the whole family $\mathcal{F}$ is smooth and one can deduce that, essentially, X_σ has normal crossing singularities. This has implications for the cohomology of $\mathrm{Trop}\, X$, at least in the constant coefficient case, which will hopefully appear in future work of either Hacking or Hacking and I. Even when $\mathcal{F}$ is not smooth, Hacking points out in a response to a question of mine that one still has a resolution of the structure sheaf of $\overline{X_0}$ by the structure sheaves of the X_σ.

Remark: One flaw of the preceding is that, if X_σ is disconnected and $\tau \supseteq \sigma$, it is possible that only part of X_σ lied in the closure of X_τ. However, there are many cases where one can exclude this possibility. In the next section, we will see that, when X is a curve, we can often deduce that the X_v are connected simply by looking at the combinatorics of $\operatorname{Trop} X$ and the degree of X.

As one easy application of what has proceeded, we prove the following results which first appeared in [6] and [17].

Proposition 2.4.5. *Suppose that X is pure (*e.g. *irreducible) of dimension d. Then* $\operatorname{Trop} X$ *is pure of dimension d.*

Proof. Let σ be a facet of Σ. Then $\bigcup_{\tau \supseteq \sigma} X_\tau = X_\sigma$. The left hand side is proper and the right hand side is affine. The only schemes that are proper and affine are the zero dimensional schemes. So we see that $\dim X_\sigma = d - \dim \sigma = 0$ and $\dim \sigma = d$. □

Proposition 2.4.6. *Suppose that X is irreducible. Then* $\operatorname{Trop} X$ *is connected.*

Proof. Suppose that $\operatorname{Trop} X = U \sqcup V$ for U and V closed and open. U and V are necessarily subcomplexes of Σ. Since these subcomplexes are closed upwards, $\bigcup_{\sigma \in U} X_\sigma$ and $\bigcup_{\sigma \in V} X_\sigma$ are closed and disjoint subsets of $\overline{X_0}$ so $\overline{X_0}$ is disconnected. But an irreducible subvariety $\overline{X}$ of a proper variety $\overline{T}$ can not degenerate within a proper family to a disconnected one, a contradiction. □

2.5 The Zero Tension Condition

Let $\Sigma \subset \mathbb{R}^n$ be a pure d-dimensional polyhedral complex, where every face has rational slope and wt a map assigning a positive integer to each facet of Σ. We call the pair (Σ, wt) a *zero tension complex* if the following condition is met: for any $(d-1)$ dimensional face ρ of Σ (a *ridge*), let $\sigma_1, \ldots, \sigma_r$ be the facets of Σ containing ρ. The image of $H(\sigma_i)$ in $\mathbb{R}^n/H(\rho)$ is a one dimensional ray with rational slope, let v_i be the minimal lattice vector along this ray. We require that, for every ρ, we have $\sum wt(\sigma_1)v_i = 0$. We will often abuse notation and refer to Σ by itself as a zero tension complex.

Theorem 2.5.1. *Recall that, if σ is a facet of* $\operatorname{Trop} X$*, then* $X_0(\sigma) \cong (\kappa^*)^d \times A$ *for some zero dimensional scheme A. Set $wt(\sigma)$ to be* $\dim_\kappa \mathcal{O}(A)$*. With this choice of wt, the complex* $\operatorname{Trop} X$ *is a zero tension complex.*

Proof. Let ρ be a codimension one face of Σ. Then, by Proposition 2.2.4, $X_\rho = C$ for some one-dimensional scheme $C \subset (\kappa^*)^n/\exp(H(\rho))$. By Proposition 2.2.3, the link of ρ is $\operatorname{Trop} C$ and one can check that the weight functions wt on $\operatorname{Trop} X$ and $\operatorname{Trop} C$ are consistent. Thus, we are reduced to the case of $\operatorname{Trop} C$ for C a curve in $(\kappa^*)^n$.

In this case, $\operatorname{Trop} C$ is a union of finitely many rays, let $v_1, \ldots, v_r$ be the minimal lattice vectors along each of these rays and let $w_i = wt(\mathbb{R}_+ v_i)$. We want to show that $\sum w_i v_i = 0$; it is enough to show that $\sum w_i < v_i, u >= 0$ for every $u \in \mathbb{Z}^n$. Each ray of $\operatorname{Trop} C$ corresponds to finitely many points of $\overline{C}$.

Consider the function $\phi : (\kappa^*)^n \to \kappa^*$ given by $(x_1, \ldots, x_n) \mapsto x_1^{u_1} \cdots x_n^{u_n}$. ϕ extends to a meromorphic function on the toric variety $\overline{T}_0$ with a pole of order $< u, v_i >$ at the boundary component of $\overline{T_0}$ assosciated to v_i. It is easy to check that ϕ extends to a meromorphic function on $\overline{C}$ with a pole of order $< u, v_i >$ at p_i. The intersection of $\overline{C}$ with this boundary component is the zero dimensional scheme whose length is defined to be w_i. So, at the points of C_σ, ϕ has $w_i < u, v_i >$ zeroes. Since ϕ has equally many zero and poles, $\sum w_i < v_i, u >= 0$. ☐

We can now explain the remark in the previous section that we can often rule out the possibility of the X_σ being disconnected by examining the geometry of $\operatorname{Trop} X$.

Proposition 2.5.2. *Suppose that $X \subset (K^*)^n$ is a curve and that, at every vertex of* $\operatorname{Trop} X$*, the edges incident to that vertex have a unique linear relation between them. Suppose furthermore that the number of unbounded rays of* $\operatorname{Trop} X$ *in direction u is equal to the degree of the function $\prod x_i^{u_i}$ on X. Then, for every vertex v of* $\operatorname{Trop} X$*, X_v is irreducible.*

Proof. Suppose for the sake of contradiction that $X_v = C_1 \cup C_2$. Then $\operatorname{Trop} X = \operatorname{Trop} C_1 \cup \operatorname{Trop} C_2$ and the weights w arising from X, C_1 and C_2 obey $w_X = w_{C_1} + w_{C_2}$. With the stated hypotheses, the only way to partition w into contributions coming from C_1 and C_2 such that the zero tension condition is obeyed is to have $\operatorname{Trop} C_1 = \operatorname{Trop} C_2$. But then $wt(e) > 1$ for each e. In particular, the unbounded rays of $\operatorname{Trop} X$ have weight greater than 1. Let u be a particular direction of unbounded ray. Then X_u must either consist of more than one point, or must consist of a point with multiplicity greater than 1. In either case, this contributes multiple zeroes to the degree of $\prod x_i^{u_i}$ and hence contradicts our assumption that $\prod x_i^{u_i}$ only has as many zeroes as there are rays in direction u. $\square$

Chapter 3

The Tropical Grassmannian

3.1 Introduction

In this chapter, we will investigate the tropicalization of the Grassmannian in the standard Plücker embedding. Explicitly, let $K[p]$ be the polynomial ring in $\binom{n}{d}$ variables indexed by the d-element subsets of $[n]$; we write the variables as $p_{i_1 \ldots i_d}$ for $1 \leq i_1 < \cdots < i_d \leq n$ and adopt the standard conventions that $p_{i_1 \cdots i_d} = (-1)^\sigma p_{i_{\sigma(1)} \cdots i_{\sigma(d)}}$ if $\sigma \in S_d$ is the permutation such that $i_{\sigma(j)}$ is increasing in j and that $p_{i_1 \cdots i_d} = 0$ if $(i_1, \ldots, i_d)$ has a repeated index.

For simplicity, assume in this chapter and the next that K and κ have the same characteristic. This implies that κ embeds in K. As all of the equations defining $G(d, n)$ have coefficients in $\mathbb{Z}$, and hence in κ, we will thus be in the contant coefficients case.

The *Plücker ideal* $I_{d,n}$ is the homogeneous prime ideal in $K[p]$ consisting of the algebraic relations among the $d \times d$-subdeterminants of any $d \times n$-matrix with entries in any commutative ring. The projective variety of $I_{d,n}$ is the *Grassmannian* $G_{d,n}$ which parameterizes all d-dimensional linear subspaces of an n-dimensional vector space.

The *tropical Grassmannian* $\mathcal{G}_{d,n}$ is $\operatorname{Trop}\operatorname{Spec} K[p^{\pm}]/I_{d,n}$. It is well known that $G_{d,n}$ has dimension $d(n-d)$ so $\operatorname{Spec} K[p^{\pm}]/I_{d,n}$ has dimension $d(n-d)+1$ and we have:

Corollary 3.1.1. *The tropical Grassmannian $\mathcal{G}_{d,n}$ is a polyhedral fan in $\mathbb{R}^{\binom{n}{d}}$. Each of its maximal cones has the same dimension, namely, $(n-d)d+1$.*

We show in Section 4.5 that $\mathcal{G}_{d,n}$ depends on the characteristic of K if $d = 3$ and $n \geq 7$. All results in this chapter are valid over any field K.

It is convenient to reduce the dimension of the tropical Grassmannian. This can be done in three possible ways. Let ϕ denote the linear map from $\mathbb{R}^n$ into $\mathbb{R}^{\binom{n}{d}}$ which sends an n-vector $(a_1, a_2, \ldots, a_n)$ to the $\binom{n}{d}$-vector whose $(i_1, \ldots, i_d)$-coordinate is $a_{i_1} + \cdots + a_{i_d}$. The map ϕ is injective, and its image is the common intersection of all cones in the tropical Grassmannian $\mathcal{G}_{d,n}$. Note that the vector $(1, \ldots, 1)$ of length $\binom{n}{d}$ lies in $\operatorname{Image}(\phi)$. We conclude:

- The image of $\mathcal{G}_{d,n}$ in $\mathbb{R}^{\binom{n}{d}}/\mathbb{R}(1, \ldots, 1)$ is a fan $\mathcal{G}'_{d,n}$ of dimension $d(n-d)$.
- The image of $\mathcal{G}_{d,n}$ or $\mathcal{G}'_{d,n}$ in $\mathbb{R}^{\binom{n}{d}}/\operatorname{Image}(\phi)$ is a fan $\mathcal{G}''_{d,n}$ of dimension $(d-1)(n-d-1)$. No cone in this fan contains a non-zero linear space.

- Intersecting $\mathcal{G}''_{d,n}$ with the unit sphere yields a polyhedral complex $\mathcal{G}'''_{d,n}$. Each maximal face of $\mathcal{G}'''_{d,n}$ is a polytope of dimension $nd - n - d^2$.

We shall distinguish the four objects $\mathcal{G}_{d,n}$, $\mathcal{G}'_{d,n}$, $\mathcal{G}''_{d,n}$ and $\mathcal{G}'''_{d,n}$ when stating our theorems below. In subsequent sections less precision is needed, and we sometimes identify $\mathcal{G}_{d,n}$, $\mathcal{G}'_{d,n}$, $\mathcal{G}''_{d,n}$ and $\mathcal{G}'''_{d,n}$ if there is no danger of confusion.

Example 3.1.2. ($d = 2, n = 4$) The smallest non-zero Plücker ideal is the principal ideal $I_{2,4} = \langle p_{12}p_{34} - p_{13}p_{24} + p_{14}p_{23}\rangle$. Here $\mathcal{G}_{2,4}$ is a fan with three five-dimensional cones $\mathbb{R}^4 \times \mathbb{R}_{\geq 0}$ glued along $\mathbb{R}^4 = \text{Image}(\phi)$. The fan $\mathcal{G}''_{2,4}$ consists of three half rays emanating from the origin (the picture of a tropical line). The zero-dimensional simplicial complex $\mathcal{G}'''_{2,4}$ consists of *three points.*

Example 3.1.3. ($d = 2, n = 5$) The tropical Grassmannian $\mathcal{G}'''_{2,5}$ is the *Petersen graph* with 10 vertices and 15 edges. This was shown in [42, Example 9.10].

The following theorem generalizes both of these examples. It concerns the case $d = 2$, that is, the tropical Grassmannian of lines in $(n-1)$-space.

Theorem 3.1.4. *The tropical Grassmannian $\mathcal{G}'''_{2,n}$ is a simplicial complex known as space of phylogenetic trees. It has $2^{n-1} - n - 1$ vertices, $1 \cdot 3 \cdots (2n{-}5)$ facets, and its homotopy type is a bouquet of $(n{-}2)\,!$ spheres of dimension $n{-}4$.*

A detailed description of $\mathcal{G}_{2,n}$ and the proof of this theorem will be given in Section 3.2. Metric properties of the *space of phylogenetic trees* were studied by Billera, Holmes and Vogtmann in [7] (our n corresponds to Billera, Holmes and Vogtmann's $n+1$.) The abstract simplicial complex and its homotopy type had been found earlier by Vogtmann [48] and by Robinson and Whitehouse [33]. The description has the following corollary. Recall that a simplicial complex is a *flag complex* if the minimal non-faces are pairs of vertices. This property is crucial for the existence of unique geodesics in [7].

Corollary 3.1.5. *The simplicial complex $\mathcal{G}'''_{2,n}$ is a flag complex.*

We do not have a complete description of the tropical Grassmannian when $d \geq 3$ and $n-d \geq 3$. We did succeed, however, in computing $\mathcal{G}_{3,6}$.

Theorem 3.1.6. *The tropical Grassmannian $\mathcal{G}'''_{3,6}$ is a 3-dimensional polyhedral complex with 65 vertices, 535 edges, 1350 triangles, 990 tetrahedra and 15 bipyramids.*

The proof and complete description of $\mathcal{G}_{3,6}$ will be presented in Section 3.3.

If L is a d-dimensional linear subspace of the vector space K^n, then $\operatorname{Trop} L$ is a polyhedral complex in $\mathbb{R}^n$. Such a polyhedral complex arising from a d-plane in K^n is called a *realizable tropical d-plane in n-space.* Since L is invariant under scaling, $\operatorname{Trop} L$ is invariant under translation by $(1,1,\ldots,1)$, so we can identify $\operatorname{Trop} L$ with its image in $\mathbb{R}^n/\mathbb{R}(1,1,\ldots,1) \simeq \mathbb{R}^{n-1}$. Thus $\operatorname{Trop} L$ becomes a $(d-1)$-dimensional polyhedral complex in $\mathbb{R}^{n-1}$. For $d=2$, we get a tree.

There is a canonical bijection between $G_{d,n}$ and the set of d-planes through the origin in K^n. The analogous bijection for the tropical Grassmannian $\mathcal{G}'_{d,n}$ is the content of the next theorem.

Theorem 3.1.7. *The bijection between the classical Grassmannian $G_{d,n}$ and the set of d-planes in K^n induces a unique bijection $w \mapsto L_w$ between the tropical Grassmannian $\mathcal{G}'_{d,n}$ and the set of tropical d-planes in n-space.*

Theorems 3.1.4, 3.1.6 and 3.1.7 are proved in Sections 3.2, 3.3 and 4.5 respectively. Almost all of the material in this chapter appeared previously in [39].

3.2 The Space of Phylogenetic Trees

In this section we prove Theorem 3.1.4 which asserts that the *tropical Grassmannian of lines* $\mathcal{G}_{2,n}$ coincides with the *space of phylogenetic trees* [7]. We begin by reviewing the simplicial complex $\mathbf{T}_n$ underlying this space.

The vertex set $\mathrm{Vert}(\mathbf{T}_n)$ consists of all unordered pairs $\{A, B\}$, where A and B are disjoint subsets of $[n] := \{1, 2, \ldots, n\}$ having cardinality at least two, and $A \sqcup B = [n]$. The cardinality of $\mathrm{Vert}(\mathbf{T}_n)$ is $2^{n-1} - n - 1$. Two vertices $\{A, B\}$ and $\{A', B'\}$ are connected by an edge in $\mathbf{T}_n$ if and only if

$$A \subset A' \quad \text{or} \quad A \subset B' \quad \text{or} \quad B \subset A' \quad \text{or} \quad B \subset B'. \tag{3.1}$$

We now define $\mathbf{T}_n$ as the flag complex with this graph. In other words, a subset

$\sigma \subseteq \mathrm{Vert}(\mathbf{T}_n)$ is a face of $\mathbf{T}_n$ if any pair $\big\{\{A, B\}, \{A', B'\}\big\} \subseteq \sigma$ satisfies (3.1).

The simplicial complex $\mathbf{T}_n$ was first introduced by Buneman (see [5, §5.1.4]) and was studied more recently by Robinson-Whitehouse [33] and Vogtmann [48]. These authors obtained the following results. Each face σ of $\mathbf{T}_n$ corresponds to a semi-labeled tree with leaves $1, 2, \ldots, n$. Here each internal node is unlabeled and has at least three neighbors. Each internal edge of such a tree defines a partition $\{A, B\}$ of the set of leaves $\{1, 2, \ldots, n\}$, and we encode the tree by the set of partitions representing its internal edges. The facets (= maximal faces) of $\mathbf{T}_n$ correspond to *trivalent trees*, that is, semi-labeled trees whose internal nodes all have three neighbors. All facets of $\mathbf{T}_n$ have the same cardinality $n-3$, the number of internal edges of any trivalent tree. Hence $\mathbf{T}_n$ is pure of dimension $n-4$. The number of facets (i.e. trivalent semi-labeled trees on $\{1, 2, \ldots, n\}$) is the *Schröder number*

$$(2n-5)!! \quad = \quad (2n-5) \times (2n-7) \times \cdots \times 5 \times 3 \times 1. \tag{3.2}$$

It is proved in [33] and [48] that $\mathbf{T}_n$ has the homotopy type of a bouquet of $(n-2)\,!$ spheres of dimension $n-4$. The two smallest cases $n = 4$ and $n = 5$ are discussed in Examples 3.1.2 and 3.1.3. Here is a description of the next case:

Example 3.2.1. $(n = 6)$ The two-dimensional simplicial complex $\mathbf{T}_6$ has 25 vertices, 105 edges and 105 triangles, each coming in two symmetry classes:

$$15 \text{ vertices like } \{12, 3456\}\,, \quad 10 \text{ vertices like } \{123, 456\},$$

$$60 \text{ edges like } \{\{12, 3456\}, \{123, 456\}\},$$

$$45 \text{ edges like } \{\{12, 3456\}, \{1234, 56\}\},$$

$$90 \text{ triangles like } \{\{12, 3456\}, \{123, 456\}, \{1234, 56\}\},$$

$$15 \text{ triangles like } \{\{12, 3456\}\}, \{34, 1256\}\}, \{56, 1234\}\}.$$

Each edge lies in three triangles, corresponding to restructuring subtrees. $\square$

We next describe an embedding of $\mathbf{T}_n$ as a simplicial fan into the $\frac{1}{2}n(n-3)$-dimensional vector space $\mathbb{R}^{\binom{n}{2}}/\text{image}(\phi)$. For each trivalent tree σ we first define a cone B_σ in $\mathbb{R}^{\binom{n}{2}}$ as follows. By a *realization* of a semi-labeled tree σ we mean a one-dimensional cell complex in some Euclidean space whose underlying graph is a tree isomorphic to σ. Such a realization of σ is a metric space on $\{1, 2, \ldots, n\}$. The *distance* between i and j is the length of the unique path between leaf i and leaf j in that realization. Then we set

$$\begin{aligned} B_\sigma \quad = \quad & \Big\{ (w_{12}, w_{13}, \ldots, w_{n-1,n}) \in \mathbb{R}^{\binom{n}{2}} \;:\; -w_{ij} \;\text{ is the distance from} \\ & \qquad \text{leaf } i \text{ to leaf } j \text{ in some realization of } \sigma \Big\} \;+\; \text{image}(\phi). \end{aligned}$$

Let C_σ denote the image of B_σ in the quotient space $\mathbb{R}^{\binom{n}{2}}/\text{image}(\phi)$. Passing to this quotient has the geometric meaning that two trees are identified if their only difference is in the lengths of the n edges adjacent to the leaves.

Theorem 3.2.2. *The closure $\overline{C}_\sigma$ is a simplicial cone of dimension $|\sigma|$ with relative interior C_σ. The collection of all cones $\overline{C}_\sigma$, as σ runs over $\mathbf{T}_n$, is a simplicial fan. It is isometric to the Billera-Holmes-Vogtmann space of trees.*

Proof. Realizations of semi-labeled trees are characterized by the *four point condition* (e.g. [5, Theorem 2.1], [9]). This condition states that for any quadruple of leaves i, j, k, l there exists a unique relabeling such that

$$w_{ij} + w_{kl} \;=\; w_{ik} + w_{jl} \;\leq\; w_{il} + w_{jk}. \tag{3.3}$$

Given any tree σ, this gives a system of $\binom{n}{4}$ linear equations and $\binom{n}{4}$ linear inequalities. The solution set of this linear system is precisely the closure $\overline{B}_\sigma$ of the cone B_σ in $\mathbb{R}^{\binom{n}{2}}$. This follows from the *Additive Linkage Algorithm* [9] which reconstructs the combinatorial tree σ from any point w in B_σ.

All of our cones share a common linear subspace, namely Image(ϕ). This is seen by replacing the inequalities in (3.3) by equalities. The cone $\overline{B}_\sigma$ is the direct sum (3.4) of this linear space with a $|\sigma|$-dimensional simplicial cone. Let $\{e_{ij} : 1 \leq i < j \leq n\}$ denote the standard basis of $\mathbb{R}^{\binom{n}{2}}$. Adopting the convention $e_{ji} = e_{ij}$, for any partition $\{A, B\}$ of $\{1, 2, \ldots, n\}$ we define

$$E_{A,B} \quad = \quad \sum_{i \in A} \sum_{j \in B} e_{ij}.$$

These vectors give the generators of our cone as follows:

$$\overline{B}_\sigma \quad = \quad \text{image}(\phi) \quad + \quad \mathbb{R}_{\geq 0} \left\{ E_{A,B} \,:\, \{A, B\} \in \sigma \right\}. \tag{3.4}$$

From the two presentations (3.3) and (3.4) it follows that

$$\overline{B}_\sigma \ \cap \ \overline{B}_\tau \quad = \quad \overline{B}_{\sigma\cap\tau} \qquad \text{for all } \sigma, \tau \in \mathbf{T}_n. \tag{3.5}$$

Therefore the cones B_σ form a fan in $\mathbb{R}^{\binom{n}{2}}$, and this fan has face poset $\mathbf{T}_n$. It follows from (3.4) that the quotient $\overline{C}_\sigma = \overline{B}_\sigma/\text{image}(\phi)$ is a pointed cone.

We get the desired conclusion for the cones $\overline{C}_\sigma$ by taking quotients modulo the common linear subspace Image(ϕ). The resulting fan in $\mathbb{R}^{\binom{n}{2}}/\text{image}(\phi)$ is simplicial of pure dimension $n-3$ and has face poset $\mathbf{T}_n$. It is isometric to the Billera-Holmes-Vogtmann space in [7] because their metric is flat on each cone $\overline{C_\sigma} \simeq \mathbb{R}^{|\sigma|}_{\geq 0}$ and extended by the gluing relations $\overline{C}_\sigma \ \cap \ \overline{C}_\tau \ = \ \overline{C}_{\sigma\cap\tau}$. □

We now turn to the tropical Grassmannian and prove our first main result. We shall identify the simplicial complex $\mathbf{T}_n$ with the fan in Theorem 3.2.2.

Proof of Theorem 3.1.4: The Plücker ideal $I_{2,n}$ is generated by the $\binom{n}{4}$ quadrics

$$p_{ij}p_{kl} \ - \ p_{ik}p_{jl} \ + \ p_{il}p_{jk} \qquad \text{for } 1 \leq i < j < k < l \leq n.$$

The tropicalization of this polynomial is the disjunction of linear systems

$$\begin{aligned} & w_{ij} + w_{kl} \ = \ w_{ik} + w_{jl} \ \leq \ w_{il} + w_{jk} \\ \text{or} \quad & w_{ij} + w_{kl} \ = \ w_{il} + w_{jk} \ \leq \ w_{ik} + w_{jl} \\ \text{or} \quad & w_{ik} + w_{jl} \ = \ w_{il} + w_{jk} \ \leq \ w_{ij} + w_{kl}. \end{aligned}$$

Every point w on the tropical Grassmannian $\mathcal{G}_{2,n}$ satisfies this for all quadruples i, j, k, l, that is, it satisfies the four point condition (3.3). The Additive Linkage

Algorithm reconstructs the unique semi-labeled tree σ with $w \in C_\sigma$. This proves that every relatively open cone of $\mathcal{G}_{2,n}$ lies in the relative interior of a unique cone C_σ of the fan $\mathbf{T}_n$ in Theorem 3.2.2.

We need to prove that the fans $\mathbf{T}_n$ and $\mathcal{G}_{2,n}$ are equal. Equivalently, every cone C_σ is actually a cone in $\operatorname{Trop} G_{d,n}$. This will be accomplished by analyzing the corresponding initial ideal. As $\operatorname{Trop} G_{d,n}$ is closed, it suffices to consider maximal faces σ of $\mathbf{T}_n$. Fix a trivalent tree σ and a weight vector $w \in C_\sigma$. Then, for every quadruple i, j, k, l, the inequality in (3.3) is strict. This means combinatorially that $\big\{\{i,l\},\{j,k\}\big\}$ is a four-leaf subtree of σ.

Let J_σ denote the ideal generated by the quadratic binomials $p_{ij}p_{kl} - p_{ik}p_{jl}$ corresponding to all four-leaf subtrees of σ. Our discussion shows that $J_\sigma \subseteq \mathrm{in}_w(I_{2,n})$. The proof will be complete by showing that the two ideals agree:

$$J_\sigma \;=\; \mathrm{in}_w(I_{2,n}). \tag{3.6}$$

This identity will be proved by showing that the two ideals have a common initial monomial ideal, generated by square-free quadratic monomials.

We may assume, without loss of generality, that $-w$ is a strictly positive vector, corresponding to a planar realization of the tree σ in which the leaves $1, 2, \ldots, n$ are arranged in circular order to form a convex n-gon (Figure 1).

Let M be the ideal generated by the monomials $p_{ik}p_{jl}$ for $1 \le i < j < k < l \le n$. These are the crossing pairs of edges in the n-gon. By a classical construction of invariant theory, known as *Kempe's circular straightening law*

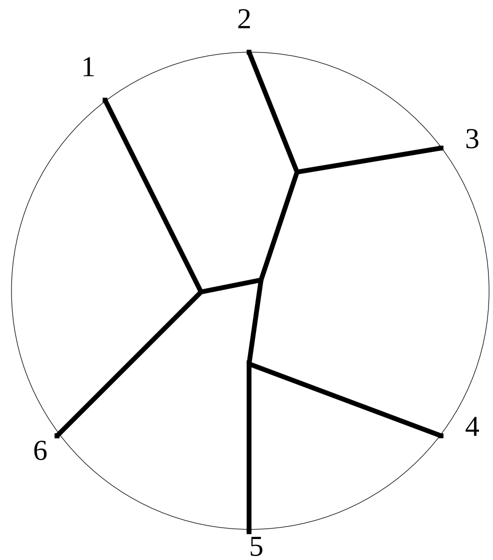

Figure 3.1: A Circular Labeling of a Tree with Six Leaves

(see [40, Theorem 3.7.3]), there exists a term order $\prec_{\text{circ}}$ on $\mathbb{Z}[p]$ such that

$$M \quad = \quad \text{in}_{\prec_{\text{circ}}}(I_{2,n}). \tag{3.7}$$

Now, by our circular choice w of realization of the tree σ, the crossing monomials $p_{ik}p_{jl}$ appear as terms in the binomial generators of J_σ. Moreover, the term order $\prec_{\text{circ}}$ on $\mathbb{Z}[p]$ refines the weight vector w. This implies

$$\text{in}_{\prec_{\text{circ}}}(\text{in}_w(I_{2,n})) \quad = \quad \text{in}_{\prec_{\text{circ}}}(I_{2,n}) \quad = \quad M \quad \subseteq \quad \text{in}_{\prec_{\text{circ}}}(J_\sigma). \tag{3.8}$$

Using $J_\sigma \subseteq \text{in}_w(I_{2,n})$ we conclude that equality holds in (3.8) and in (3.6). □

3.3 The Grassmannian of 3-planes in 6-space

In this section we study the case $d = 3$ and $n = 6$. The Plücker ideal $I_{3,6}$ is minimally generated by 35 quadrics in the polynomial ring in 20 variables,

$$\mathbb{Z}[p] \quad = \quad \mathbb{Z}[p_{123}, p_{124}, \ldots, p_{456}].$$

We are interested in the 10-dimensional fan $\mathcal{G}_{3,6}$ which consists of all vectors $w \in \mathbb{R}^{20}$ such that $\mathrm{in}_w(I_{3,6})$ is monomial-free. The four-dimensional quotient fan $\mathcal{G}''_{3,6}$ sits in $\mathbb{R}^{20}/\mathrm{image}(\phi) \simeq \mathbb{R}^{14}$ and is a cone over the three-dimensional polyhedral complex $\mathcal{G}'''_{3,6}$. Our aim is to prove Theorem 3.1.6, which states that $\mathcal{G}'''_{3,6}$ consists of 65 vertices, 535 edges, 1350 triangles, 990 tetrahedra and 15 bipyramids.

We begin by listing the vertices. Let E denote the set of 20 standard basis vectors E_{ijk} in $\mathbb{R}^{\binom{6}{3}}$. For each 4-element subset $\{i, j, k, l\}$ of $\{1, 2, \ldots, 6\}$ we set

$$F_{ijkl} \quad = \quad E_{ijk} + E_{ijl} + E_{ikl} + E_{jkl}.$$

Let F denote the set of these 15 vectors. Finally consider any of the 15 *tripartitions* $\{\{i_1, i_2\}, \{i_3, i_4\}, \{i_5, i_6\}\}$ of $\{1, 2, \ldots, 6\}$ and define the vectors

$$G_{i_1 i_2 i_3 i_4 i_5 i_6} := F_{i_1 i_2 i_3 i_4} + E_{i_3 i_4 i_5} + E_{i_3 i_4 i_6}$$

$$\text{and} \qquad G_{i_1 i_2 i_5 i_6 i_3 i_4} := F_{i_1 i_2 i_5 i_6} + E_{i_3 i_5 i_6} + E_{i_4 i_5 i_6}.$$

This gives us another set G of 30 vectors. All 65 vectors in $E \cup F \cup G$ are regarded as elements of the quotient space $\mathbb{R}^{\binom{6}{3}}/\mathrm{image}(\phi) \simeq \mathbb{R}^{14}$. Note that

$$G_{i_1 i_2 i_3 i_4 i_5 i_6} \quad = \quad G_{i_3 i_4 i_5 i_6 i_1 i_2} \quad = \quad G_{i_5 i_6 i_1 i_2 i_3 i_4}.$$

Later on, the following identity will turn out to be important in the proof of Theorem 3.3.4:

$$G_{i_1 i_2 i_3 i_4 i_5 i_6} + G_{i_1 i_2 i_5 i_6 i_3 i_4} \quad = \quad F_{i_1 i_2 i_3 i_4} + F_{i_1 i_2 i_5 i_6} + F_{i_3 i_4 i_5 i_6}. \tag{3.9}$$

Lemma 3.3.1 and other results in this section were found by computation.

Lemma 3.3.1. *The set of vertices of* $\mathcal{G}_{3,6}$ *equals* $E \cup F \cup G$.

We next describe all the 550 edges of the tropical Grassmannian $\mathcal{G}_{3,6}$.

(EE) There are 90 edges like $\{E_{123}, E_{145}\}$ and 10 edges like $\{E_{123}, E_{456}\}$, for a total of 100 edges connecting pairs of vertices both of which are in E. (By the word "like", we will always mean "in the S_6 orbit of, where S_6 permutes the indices $\{1, 2, \ldots 6\}$.")

(FF) This class consists of 45 edges like $\{F_{1234}, F_{1256}\}$.

(GG) Each of the 15 tripartitions gives exactly one edge, like $\{G_{123456}, G_{125634}\}$.

(EF) There are 60 edges like $\{E_{123}, F_{1234}\}$ and 60 edges like $\{E_{123}, F_{1456}\}$, for a total of 120 edges connecting a vertex in E to a vertex in F.

(EG) This class consists of 180 edges like $\{E_{123}, G_{123456}\}$. The intersections of the index triple of the E vertex with the three index pairs of the G vertex must have cardinalities $(2, 1, 0)$ *in this cyclic order*.

(FG) This class consists of 90 edges like $\{F_{1234}, G_{123456}\}$.

Lemma 3.3.2. *The* 1*-skeleton of* $\mathcal{G}'''_{3,6}$ *is the graph with the* 550 *edges above.*

Let Δ denote the *flag complex* specified by the graph in the previous lemma. Thus Δ is the simplicial complex on $E \cup F \cup G$ whose faces are subsets σ with the property that each 2-element subset of σ is one of the 550 edges. We will see that $\mathcal{G}_{3,6}$ is a subcomplex homotopy equivalent to Δ.

Lemma 3.3.3. *The flag complex* Δ *has* $1,410$ *triangles,* $1,065$ *tetrahedra,* 15 *four-dimensional simplices, and it has no faces of dimension five or more.*

The facets of Δ are grouped into seven symmetry classes:

Facet FFFGG: There are 15 four-dimensional simplices, one for each partition of $\{1,\ldots,6\}$ into three pairs. The tripartition $\{\{1,2\},\{3,4\},\{5,6\}\}$. gives the facet $\{F_{1234}, F_{1256}, F_{3456}, G_{123456}, G_{125634}\}$; the other 14 cases can be determined by the S_6 symmetry. The 75 tetrahedra contained in these 15 four-simplices are not facets of Δ.

The remaining 990 tetrahedra in Δ are facets and they come in six classes:

Facet EEEE: There are 30 tetrahedra like $\{E_{123}, E_{145}, E_{246}, E_{356}\}$.

Facet EEFF1: There are 90 tetrahedra like $\{E_{123}, E_{456}, F_{1234}, F_{3456}\}$.

Facet EEFF2: There are 90 tetrahedra like $\{E_{125}, E_{345}, F_{3456}, F_{1256}\}$.

Facet EFFG: There are 180 tetrahedra like $\{E_{345}, F_{1256}, F_{3456}, G_{123456}\}$.

Facet EEEG: There are 240 tetrahedra like $\{E_{126}, E_{134}, E_{356}, G_{125634}\}$.

Facet EEFG: There are 360 tetrahedra like $\{E_{234}, E_{125}, F_{1256}, G_{125634}\}$.

While Δ is an abstract simplicial complex on the vertices of $\mathcal{G}'''_{3,6}$, it is not embedded as a simplicial complex because relation (3.9) shows that the five vertices of the four dimensional simplices only span three dimensional space. Specifically, they form a bipyramid with the F-vertices as the base and the G-vertices as the two cone points.

We now modify the flag complex Δ to a new polyhedral complex Δ' which has pure dimension three and reflects the situation described in the last paragraph. The complex Δ' is obtained from Δ by removing the 15 FFF-triangles $\{F_{1234}, F_{1256}, F_{3456}\}$, along with the 30 tetrahedra FFFG and the 15 four-dimensional facets FFFGG containing the FFF-triangles. In the place of each four dimensional simplex, we instead put a bipyramid. We will give another way of understanding the seven types of facets and three types of rays in Section 4.3.

The following theorem implies Theorem 3.1.6.

Theorem 3.3.4. *The tropical Grassmannian $\mathcal{G}'''_{3,6}$ equals the polyhedral complex Δ'. It is not a flag complex because of the 15 missing FFF-triangles. The homology of $\mathcal{G}'''_{3,6}$ is concentrated in (top) dimension 3; $H_3(\mathcal{G}'''_{3,6}, \mathbb{Z}) = \mathbb{Z}^{126}$.*

The integral homology groups were computed independently by Michael Joswig and Volkmar Welker. We are grateful for their help.

This theorem is proved by an explicit computation. The correctness of the result can be verified by the following method. One first checks that the seven types of cones described above are indeed Gröbner cones of $I_{3,6}$ whose initial ideals are monomial-free. Next one checks that the list is complete. This relies on a result that will appear in a forthcoming paper which guarantees that $\mathcal{G}_{3,6}$ is connected in codimension 1. The completeness check is done by computing the link of each of the known classes of triangles. Algebraically, this amounts to computing the (truly zero-dimensional) tropical variety of $\mathrm{in}_w(I_{3,6})$ where w is any point in the relative interior of the triangular cone in question. For all but one class of triangles the link consists of three points, and each neighboring 3-cell is found to be already among our seven classes. The links of the triangles are as follows:

Triangle EEE: The link of $\{E_{146}, E_{256}, E_{345}\}$ consists of E_{123}, G_{163425}, G_{142635}.

Triangle EEF: The link of $\{E_{256}, E_{346}, F_{1346}\}$ consists of F_{1256}, G_{132546}, G_{142536}.

Triangle EEG: The link of $\{E_{156}, E_{236}, G_{142356}\}$ consists of E_{124}, E_{134}, F_{1456}.

Triangle EFF: The link of $\{E_{135}, F_{1345}, F_{2346}\}$ consists of E_{236}, E_{246}, G_{153426}.

Triangle EFG: The link of $\{E_{235}, F_{2356}, G_{143526}\}$ consists of E_{145}, F_{1246}, E_{134}.

Triangle FGG: The link of $\{F_{1456}, G_{142356}, G_{145623}\}$ consists of F_{2356} and F_{1234}.

Chapter 4

Tropical Linear Spaces

4.1 Introduction and Summary

In the preceding section we have described our attempts to study the tropicalization of the Grassmannian. Classically, the Grassmannian parameterizes linear spaces of dimension d in n space. A similar statement is true of the Tropical Grassmannian, namely:

Theorem 4.1.1. *There is a bijection between the points of* $\operatorname{Trop} G(d, n)$ *and the collection of polyhedral complexes occurring as* $\operatorname{Trop} L$ *for* L *a* d*-plane in* K^n*.*

Just as it is best to study all matroids and then consider the realizable matroids as a special subset among them, it turns out to be best to enhance the notion of "$\operatorname{Trop} L$ for L a d-plane in K^n" to a larger collection of combinatorial objects which we will term tropical linear spaces. After we have investigated tropical linear spaces from a combinatorial perspective, we then turn in Section 4.5 to the question of identifying which such spaces occur as $\mathrm{Trop}\, L$. The reader who prefers combinatorics to algebraic geometry may find this chapter a pleasant break, as most of our arguments will be purely combinatorial.

Tropical linear spaces will be polyhedral complexes. Thus, we can ask about all of the invariants associated to polyhedral complexes. Ardila and Klivans [1] have studied the link of a vertex in a tropical linear space and have shown it to be homeomorphic to the chain complex of the lattice of flats of a certain matroid and thus, ([50], sect. 7.6), homotopic to a wedge of spheres. The polytopes occurring in tropical linear spaces are Minkowski summands of permutahedra (Proposition 4.2.5). The topology of tropical linear spaces is quite simple: they are contractible (Theorem 4.2.8). So we understand the local combinatorics and the local and global topology of tropical linear spaces.

The global combinatorics of tropical linear spaces, on the other hand, is quite intriguing. There is a great deal of theoretical and some experimental evidence for the following:

The f-Vector Conjecture. The number of i-dimensional faces of a tropical d-plane in n-space which become bounded after being mapped to $\mathbb{R}^n/(1,\ldots,1)$ is at most $\binom{n-2i}{d-i}\binom{n-i-1}{i-1}$.

Remark: It would follow from the f-vector conjecture that the number of total i-dimensional faces of a tropical d-plane in n-space, without any boundedness condition, is at most $\binom{n-i-1}{d-i}\binom{2n-d-1}{i-1}$. See Proposition 4.2.10.

The following is a summary of the rest of the chapter.

In Section 4.2, we will define the basic concepts mentioned in this introduction and prove their essential properties. We will also introduce an operation called dualization, which is analogous to the classical operation of taking the dual, or orthogonal complement, of a linear space.

In Section 4.4, we will introduce an operation of transverse intersection that produces new tropical linear spaces from old. We define a tropical linear space to be *constructible* if it can be built from tropical hyperplanes (equivalently, from points) by successive dualization and transverse intersection. Much of the rest of the chapter will be devoted toward proving:

Theorem 4.1.2. *Every constructible space achieves the f-vector of the f-vector conjecture.*

In Section 4.5, we discuss which tropical linear spaces are of the form $\operatorname{Trop} L$ for a linear space L or, as we will term it there, which are realizable. These section includes many counter-examples to show that tropical linear spaces can fail to be realizable in almost every conceivable way.

In Section 4.6 we will define a notion of "series-parallel tropical linear space", analogous to the notion of series-parallel matroid. We also describe a way of working with series-parallel matroids in terms of two-colored trees that will be important in the future. One of our themes will be that series-parallel tropical linear spaces are the best tropical linear spaces. In particular we conjecture

The f-Vector Conjecture, Continued. Equality in the f-vector conjecture is achieved precisely by series-parallel linear spaces.

We will show that the tropical linear spaces which are easy to write down are series parallel. More precisely,

Theorem 4.1.3. *Every constructible space is series-parallel.*

One unfortunate consequence of Theorem 4.1.3 is that it is quite hard to find a general method for writing down tropical linear spaces that are not series-parallel! More precisely, it is not so hard to produce degenerate limits of series-parallel linear spaces which are not themselves series-parallel. Writing down a large number of tropical linear spaces which are not limits of series-parallel tropical linear spaces, however, is a challenge – which is one of the reasons that experimentally testing the f-vector conjecture is tricky.

In Section 4.7, we will prove the f-vector conjecture in the cases $i = 1$ and $d = \lfloor n/2 \rfloor$. We then return to proving our main results. We prove Theorem 4.1.3 in Section 4.8 and also prove many lemmas that will be used in the proof of Theorem 4.1.2. We also prove Theorem 4.1.2 in Section 4.8.

Finally, in Section 4.9, we will introduce a notion of tree linear space which achieves the bounds in the f-vector conjecture and has very explicit combinatorics. Suggestively, the construction of tree linear spaces will be reminiscent of the construction of cyclic polytopes.

4.2 Basic Definitions and Results

We will adopt the convention of writing $\mathrm{Trop}(f)$ to mean the tropicalization of the hypersurface generated by f.

Consider a collection of variables $p_{i_1 \ldots i_d}$ indexed by the d element subsets of $[n] := \{1, \ldots, n\}$. We will say that p is a *tropical Plücker vector* if

$$(p_I) \in \mathrm{Trop}(P_{Sij}P_{Skl} - P_{Sik}P_{Sjl} + P_{Sil}P_{Sjk})$$

for every $S \in \binom{[n]}{d-2}$ and every i, j, k and l distinct members of $[n] \setminus S$.

Remark: This definition is equivalent to saying that p_I is a valuated matroid with values in the semiring $(\mathbb{R}, +, \min)$. See [13] as well as the more general [12] and [14].

Let $\Delta(d, n)$ denote the (d, n)-hypersimplex, defined as the convex hull of the points $e_{i_1} + \cdots + e_{i_d} \in \mathbb{R}^n$ where $\{i_1, \ldots, i_d\}$ runs over $\binom{[n]}{d}$. We abbreviate $e_{i_1} + \cdots + e_{i_d}$ by $e_{i_1 \ldots i_d}$. Consider a real-valued function $\{i_1, \ldots, i_d\} \mapsto p_{i_1 \ldots i_d}$ on the vertices of $\Delta(d, n)$. We define a polyhedral subdivision $\mathcal{D}_p$ of $\Delta(d, n)$ as follows: consider the points $(e_{i_1} + \cdots + e_{i_d}, p_{i_1 \cdots i_d}) \in \Delta(d, n) \times \mathbb{R}$ and take their convex hull. Take the lower faces (those whose outward normal vector have last component negative) and project them back down to $\Delta(d, n)$, this gives us the subdivision $\mathcal{D}_p$. We will often omit the subscript p when it is clear from context. A subdivision which is obtained in this manner is called *regular*, see, for example, [51] definition 5.3.

Let P be a subpolytope of $\Delta(d, n)$. We say that P is matroidal if the vertices of P, considered as elements of $\binom{[n]}{d}$, are the bases of a matroid M. In this case, we write $P = P_M$. Our references for matroid terminology and theory are Neil White's anthologies [49] and [50].

Proposition 4.2.1. *The following are equivalent:*

1. $p_{i_1 \ldots i_d}$ *are tropical Plücker coordinates*

2. *The one skeleta of $\mathcal{D}$ and $\Delta(d, n)$ are the same.*

3. *Every face of $\mathcal{D}$ is matroidal.*

Proof. (1) $\implies$ (2). Every edge of $\mathcal{D}$ joins two vertices of $\Delta(d, n)$. If e is an edge of $\mathcal{D}$ connecting e_I and e_J, we define the *length* of e, denoted $\ell(e)$, to be $|I \setminus J| = |J \setminus I|$. Our claim is that every edge of e has length 1. We prove by induction on $\ell \geq 2$ that $\mathcal{D}$ has no edge of length ℓ.

For the base case, if e is an edge with $\ell(e) = 2$ then $e = (e_{Sij}, e_{Skl})$ for some $S \in \binom{[n]}{d-2}$. The six vertices e_{Sij}, e_{Sik}, e_{Sil}, e_{Sjk}, e_{Sjl}, e_{Skl} form the vertices of an octahedron O with e_{Sij} and e_{Skl} opposite vertices. One can check that the condition $(p_I) \in \mathrm{Trop}(P_{Sij}P_{Skl} - P_{Sik}P_{Sjl} + P_{Sil}P_{Sjk})$ implies that O is either a face of $\mathcal{D}$ or subdivided in $\mathcal{D}$ into two square pyramids (in one of three possible ways). In any of these cases, e is not an edge of $\mathcal{D}$.

Now consider the case where $\ell > 2$. Suppose (for contradiction) that e is an edge of $\mathcal{D}$. Let $e = (e_{ST}, e_{ST'})$ with $T \cap T' = \emptyset$ and $|T| = |T'| = \ell$. Let F be the face of $\Delta(d, n)$ consisting of all vertices e_I with $S \subset I \subset S \cup T \cup T'$. Then

e must belong to some two dimensional face of $\mathcal{D}$ contained in F; call this two dimensional face G.

Let γ be the path from e_{ST} to $e_{ST'}$ that goes around G the other way from e. No two vertices of F are more than distance ℓ apart, so the edges of γ have lengths less than or equal to ℓ. If γ contained an edge $(e_{SU}, e_{SU'})$ of length ℓ then its midpoint $(e_{SU} + e_{SU'})/2$ would also be the midpoint of e contradicting the convexity of G. Thus, every edge of γ has length less than ℓ and by induction must have length 1. So all the edges of γ are in the direction $e_i - e_j$ for some i and $j \in [n]$. These vectors must span a two dimensional space. This means either that there are $\{i_1, i_2, j_1, j_2\} \subset [n]$ such that all edges of γ are parallel to some $e_{i_r} - e_{j_s}$ or there are $\{i_1, i_2, i_3\} \subset [n]$ such that every edge of γ is parallel to $e_{i_r} - e_{i_s}$. In either case, e has length at most 2, but that returns us to our ground case.

(2) $\implies$ (3): Let P be any polytope in $\mathcal{D}$. By assumption, all of the edges of P are edges of $\Delta(d, n)$. It is a theorem of Gelfand, Goresky, MacPherson and Serganova ([20], theorem 4.1) that this implies that P is matroidal. Since the proof is short, we include it. Let e_I and e_J be vertices of P with $j \in J \setminus I$. We must prove that there is a vertex of P of the form $e_{I \cup \{j\} \setminus \{b\}}$ for some $b \in I$.

Define a linear functional $\phi : \Delta(d, n) \to \mathbb{R}$ by $\phi(x_1, \ldots, x_n) = \sum_{i \in I} x_i + dx_j$. Then $Q := P \cap \{x : \phi(x) \geq d\}$ is a convex polytope and hence connected. Q contains the vertices e_I and e_J and there is a path from e_I to e_J along edges of P which lie in Q. Let the first step of this path go from e_I to $e_{I\cup\{a\}\setminus\{b\}}$. If $a \neq j$ then $\phi(e_{I\cup\{a\}\setminus\{b\}}) = d-1$ and $e_{I\cup\{a\}\setminus\{b\}}$ does not lie in Q. So instead we have $a = j$ and we are done.

(3) $\implies$ (1): It is easy to check that, if (1) is false, $\mathcal{D}$ has a one dimensional face of the form $\mathrm{Hull}(e_{Sij}, e_{Skl})$, with i, j, k and l distinct. This is not matroidal. $\square$

Now, suppose that $(p_I) \in \mathbb{R}^{\binom{[n]}{d}}$ obeys the tropical Plücker relations. Define $L(p) \subset \mathbb{R}^n$ by

$$\bigcap_{1\leq j_1<\cdots<j_{d+1}\leq n} \mathrm{Trop}(\sum_{r=1}^{d+1}(-1)^r P_{j_1\cdots\hat{j_r}\cdots j_{d+1}} X_{j_r}).$$

We term any L which arises in this manner a *d-dimensional tropical linear space in n-space*. We often omit the dependence on p when it is clear from context.

$L(p)$ is essentially the same set as the tight span of p defined in [15]. There are two differences: (1) Dress's sign conventions are opposite to ours and (2) $L(p)$ is invariant under translation by $(1,\ldots,1)$; Dress chooses a particular representative within each orbit for this translation.

While the above definition makes the connection to ordinary linear spaces most clear, for practically every purpose it is better to work with the alternate characterization which we now give. For any $w \in \mathbb{R}^n$, define $\mathcal{D}_w$ to be the subset of the vertices of $\Delta(d,n)$ at which $p_{i_1\ldots i_d} - \sum_{j=1}^{d} w_{i_j}$ is minimal. This is, by definition, a face of $\mathcal{D}$ and thus is P_{M_w} for a matroid M_w.

Proposition 4.2.2. *$w \in L(p)$ if and only if M_w is loop-free.*

Recall that a matroid is called loop-free if every element of the matroid appears in at least one basis. There is a geometric way of recognizing when M is loop-free from the polytope P_M: M is loop-free if and only if P_M is not contained in any of the n facets of $\Delta(d,n)$ of the type $x_i = 0$ for $1 \leq i \leq n$. In particular, note that if P_M meets the interior of $\Delta(d,n)$ then M is necessarily loop-free.

Proof. By replacing $p_{i_1 \dots i_d}$ by $p_{i_1 \dots i_d} - \sum_{j=1}^{d} w_{i_j} -$ (constant), we may assume without loss of generality that $w = 0$ and $\min p_I = 0$. Then $M_w = M_0$ is simply the matroid whose basis correspond to the I for which $p_I = 0$. First, we assume that M_0 has a loop j and prove that $0 \notin L(p)$. Let $(i_1, \dots, i_d)$ be a basis of M_0, clearly $j \notin (i_1, \dots, i_d)$. Then $p_{i_1 \dots i_d} = 0$ but $p_{j i_1 \dots \hat{i_r} \dots i_d} > 0$ for $1 \leq r \leq d$. Taking $(j_1, \dots, j_{d+1}) = (j, i_1, \dots, i_d)$ we see that $0 \notin \mathrm{Trop}(\sum_{r=1}^{d+1} (-1)^r P_{j_1 \dots \hat{j_r} \dots j_{d+1}} X_{j_r})$.

The converse is more interesting. Fix $J = \{j_1, \dots, j_{d+1}\}$, our aim will be to prove $0 \in \mathrm{Trop}(\sum_{r=1}^{d+1} (-1)^r P_{j_1 \dots \hat{j_r} \dots j_{d+1}} X_{j_r})$. Let $e_J = \sum e_{j_r}$ and, for any $s \in \mathbb{R}$, set $M_s = M_{-se_J}$. Note that, for s large enough, all of the bases of M_s will be subsets of J. It is equivalent to show that, for such an s, the matroid M_s has at least two bases. In other words, we must show that for any $j \in J$ and any such s, j is not a loop of M_s. By hypothesis, j is not a loop of M_0, so it is enough to show that if $j \in J$ is a loop of M_s for some s then it is a loop of $M_{s'}$ for all $s' < s$.

As s varies, M_s changes at a finite number of values of s, call them $s_1 < s_2 < \dots < s_k$. Suppose that $s_i < s < s_{i+1}$, then P_{M_s} is a face of both $P_{M_{s_i}}$

and of $P_{M_{s_{i+1}}}$. The bases of M_s are precisely the bases of M_{s_i} which have the largest possible number of elements in common with J. Similarly, the bases of M_s are precisely the bases of $M_{s_{i+1}}$ which have the smallest possible number of elements in common with J. In other words,

$$M_s = M_{s_i}|_J \oplus M_{s_i}/J = M_{s_{i+1}}|_{[n]\setminus J} \oplus M_{s_{i+1}}/([n] \setminus J).$$

From the displayed equation, it follows that if $j \in J$ is a loop in $M_{s_{i+1}}$ then it is a loop in M_s. Similarly, if j is a loop in M_s then it is a loop in M_{s_i}. Concatenating deductions of this sort, we see that, as promised, if j is a loop of M_s then it is a loop of $M_{s'}$ for all $s' < s$. $\square$

As $\Delta(d, n)$ is contained in the hyperplane $x_1 + \cdots + x_n = d$ we see that L is invariant under translation by $(1, 1, \ldots, 1)$. We will abuse notation by saying a face of L is bounded if its image in $\mathbb{R}^n/(1, \ldots, 1)$ is bounded and calling a face a vertex if its image in $\mathbb{R}^n/(1, \ldots, 1)$ is zero dimensional. However, when we refer to the dimensions of faces of L we will always be speaking of L itself, without taking the quotient by $(1, \ldots, 1)$.

We see that L is a subcomplex of $\mathcal{D}^\vee$, where $\mathcal{D}^\vee$ is defined to be the polyhedral subdivision of $\mathbb{R}^n$ where w and w' lie in the same face if $M_w = M_{w'}$. In Figure 4.1, we show $\Delta(2, 4)$, which is an octahedron, subdivided into two square pyramids and draw the dual $L/(1, 1, 1, 1)$ in bold. Notice that four of the triangular faces of the octahedron correspond to a loop-free matroid consisting of two parallelism classes, of sizes 1 and 3, while the other four faces correspond

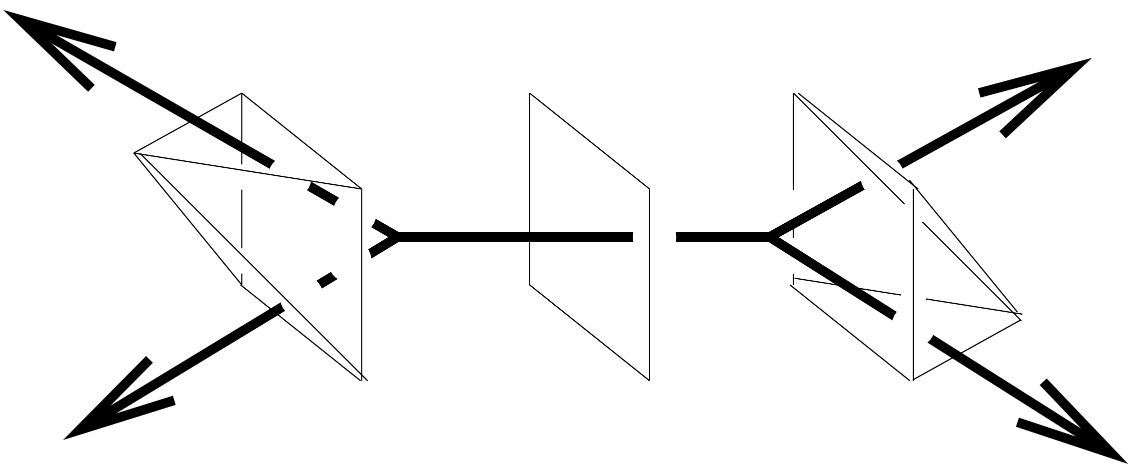

Figure 4.1: A Subdivision of $\Delta(2,4)$ and the Corresponding $L/(1,1,1,1)$

to the matroid consisting of a single loop and no other relations. We write $M^\vee$ for the face of $\mathcal{D}^\vee$ dual to P_M. If $v = M^\vee/(1,\ldots,1)$ is a vertex of $L/(1,\ldots,1)$ we see that the link of $L/(1,\ldots,1)$ at v is the subcomplex of the normal fan to P_M consisting of the normals to the loop-free facets. This fan is studied in [1], we summarize the main results of that paper in the following proposition:

Proposition 4.2.3. $\mathrm{link}_{M^\vee}(L)$ *is homeomorphic to the chain complex of the lattice of flats of L. If $N^\vee$ is a face of L containing v then $N = \bigoplus Q_{k+1}/Q_k$ for some flag of flats $\emptyset = Q_0 \subset Q_1 \subset \cdots \subset Q_r = [n]$ of M. The local cone of N at M is given by those vectors which are constant on the set $Q_k \setminus Q_{k-1}$ and for which $X(Q_k \setminus Q_{k-1}) > X(Q_{k-1} \setminus Q_{k-2})$*

When $M^\vee/(1,\ldots,1)$ is positive dimensional, P_M turns out to be a product of P_N's.

Proposition 4.2.4. *Let M be a matroid with $M = \bigoplus M_i$ the decomposition of M into connected parts. Then $P_M = \prod P_{M_i}$.*

Proof. By the definition of direct sum, the bases of M are all combinations of bases of the M_i. $\square$

We can now describe the faces of L. Recall that the braid arrangement $\mathcal{B}_r$ is the fan in $\mathbb{R}^r$ whose facets are the cones of the type $x_{\sigma(1)} \leq \ldots \leq x_{\sigma(r)}$ where σ ranges over the symmetric group S_r.

Proposition 4.2.5. *If $M^\vee$ is an r-dimensional face of L then the normal fan to $M^\vee$ is a subfan of a coarsening of $\mathcal{B}_r$. When $M^\vee$ is bounded, $M^\vee$ is a Minkowski summand of the r^{th} permutahedron.*

Proof. The normal fan to $M^\vee$ is the same as the link of P_M in $\mathcal{D}$. Let $P_{\tilde{M}}$ be a face of $\mathcal{D}$ with P_M a face of $P_{\tilde{M}}$. We must show that the local cone of $P_{\tilde{M}}$ at P_M is a union of cones of $\mathcal{B}_r$. But the local cone of $P_{\tilde{M}}$ at P_M is precisely the negative of the cone of the normal fan to $P_{\tilde{M}}$ dual to P_M. From the description in [1], we see that the normal fan to $P_{\tilde{M}}$ is a coarsening of $\mathcal{B}_n$ and that the cones dual to codimension i faces are unions of cones of $\mathcal{B}_r$.

If $M^\vee$ is bounded then its normal fan is complete and thus, by the last paragraph, a coarsening of $\mathcal{B}_r$. $\mathcal{B}_r$ is the normal fan to the r^{th} permutahedron. If P and Q are two polytopes such that the normal fan to Q refines that of P then P is a Minkowski summand of tQ for t sufficiently large. $\square$

Not every coarsening of $\mathcal{B}_r$ is the normal fan to a polytope. For example, take a cube, divide one of its faces into two triangles and cone from the center. This is a coarsening of $\mathcal{B}_4$ which is not the normal fan of any polytope. I do not know whether there are further obstructions to a fan showing up as the normal fan to a face of a tropical linear space.

Problem 4.2.6. Let P be a polytope whose normal fan is a coarsening of $\mathcal{B}_r$. Can P always occur as a face of a tropical linear space?

For ordinary linear spaces, the linear space determines the Plücker coordinates up to scaling. For tropical linear spaces, a similar result holds.

Proposition 4.2.7. *Let L be a tropical linear space with tropical Plücker coordinates p_I. Then L determines the p_I up to addition of a common scalar.*

Proof. Let S be a $d-1$ element subset of $[n]$ and let F denote the face of $\Delta(d,n)$ whose vertices are those I containing S. Since F is an $(n-d)$-simplex, it can not be subdivided in $\mathcal{D}$ so F is a face of $\mathcal{D}$. F corresponds to a loop-free matroid Φ so it is dual to a face $\Phi^\vee$ of L. $\Phi^\vee$ is of the form $(\mathbb{R}^S_{\geq 0} + \mathbb{R}(1,\ldots,1)) + q$ for q a vector satisfying $q_i = p_{Si}$ for $i \not\in S$. No other face of L is of the form $(\mathbb{R}^S_{\geq 0} \oplus \mathbb{R}(1,\ldots,1)) + q$ and q is unique up to translation by $(1,\ldots,1)$. Thus, L determines $p_{Si} - p_{Sj}$. As this is true for all S, i and j, we see that L determines the p_I up to adding the same constant to all of them. ☐

Theorem 4.2.8. *L is a pure d-dimensional contractible polyhedral complex.*

Proof. To show that L is pure d-dimensional, it is enough to show that the link of any vertex v of $L/(1,\ldots,1)$ is a pure $(d-2)$-dimensional complex. Let $v+(1,\ldots,1)$ be dual to P_M for some matroid M. By the results of [1], the link of v has a subdivision isomorphic to the order complex of the lattice of flats of M. This lattice is graded of length d, so its order complex is pure of dimension $d-2$.

We now show contractibility. It is easy to see that each tropical hyperplane is tropically convex in the sense of [11]. Thus, their intersection is tropically convex and hence contractible. □

We now study the bounded faces of L.

Proposition 4.2.9. *The following are equivalent:*

1. *w lies in a bounded face of L*
2. *P_{M_w} is an interior face of $\mathcal{D}$.*
3. *M_w is loop-free and co-loop-free.*

Proof. (2) $\iff$ (3): P_{M_w} is interior if and only if it is not contained in any facet of $\Delta(d,n)$. The facets of $\Delta(d,n)$ are of two types: $x_i = 1$ and $x_i = 0$. P_{M_w} lies in the former type if and only if i is a co-loop of M_w; P_{M_w} lies in the latter type if and only if i is a loop of M_w.

(1) $\iff$ (2): If P is not an interior face of $\mathcal{D}$, the corresponding dual face is not bounded. Conversely, if P is internal, the corresponding dual face of

$\mathcal{D}^\vee$ is bounded, so we just must check that this dual face is in $L(p)$ at all. For this, we must check that M_w is loop-free. But we saw in the previous paragraph that this follows from P_{M_w} being internal. □

At this point, we can prove our earlier remark.

Proposition 4.2.10. *If the f-vector conjecture is true then a tropical d-plane in n-space has no more than $\binom{n-i-1}{d-i}\binom{2n-d-1}{i-1}$ faces of dimension i.*

Proof. Let $\mathcal{D}$ be a matroidal decomposition of $\Delta(d,n)$. We must bound the number of codimension $i-1$ faces P_M of $\mathcal{D}$ such that M has no loops. For such a P_M, let $\Lambda \subset [n]$ be the co-loops of M and let Δ be the face of $\Delta(d,n)$ whose vertices are the sets $I \in \binom{[n]}{d}$ which contain Λ. The claim that M has no loops is equivalent to saying that P_M is an interior face of the decomposition of Δ. So we may count the total number of P_M by summing over all such faces Δ.

There are $\binom{n}{j}$ faces Δ for which $|\Lambda| = j$; each such Δ is isomorphic to $\Delta(d-j, n-j)$. A subpolytope of such a Δ has codimension $i-1$ in $\Delta(d,n)$ if and only if it has co-dimension $i-j-1$ in Δ. So, assuming the f-vector conjecture and writing the bound in the f-vector conjecture as $\binom{n-i-1}{d-i}\binom{n-d-1}{i-1}$, we have the following bound for the number of co-dimension $i-1$ loop-free faces of $\mathcal{D}$:

$$\sum_j \binom{n}{j} \binom{(n-j)-(i-j)-1}{(d-j)-(i-j)} \binom{(n-j)-(d-j)-1}{(i-j)-1} =$$
$$\binom{n-i-1}{d-i} \sum_j \binom{n}{j} \binom{n-d-1}{i-j-1} =$$
$$\binom{n-i-1}{d-i} \binom{2n-d-1}{i-1}.$$

The last sum is evaluated by the identity $\sum_j \binom{p}{j}\binom{q}{r-j} = \binom{p+q}{r}$. □

Suppose that $p : \binom{[n]}{d} \to \mathbb{R}$ obey the tropical Plücker relations. Define $p^{\perp} : \binom{[n]}{n-d} \to \mathbb{R}$ by $p^{\perp}_I = p_{[n]\setminus I}$. Then the subdivision $\mathcal{D}^{\perp}$ of $\Delta(n-d,n)$ induced by $p^{\perp}$ is simply the image of $\mathcal{D}$ under $(x_1,\ldots,x_n) \mapsto (1-x_1,\ldots,1-x_n)$. This map replaces each matroidal polytope with the polytope of the dual matroid; in particular, $p^{\perp}$ also obeys the tropical Plücker relations. We denote $L(p^{\perp})$ by $L^{\perp}$. We can easily check that $L \mapsto L^{\perp}$ is an inclusion reversing bijection from the set of tropical linear spaces in $\mathbb{R}^n$ to itself.

We will use this dualization operation to defien a large class of constructible tropical linear spaces, which are built out of repeated simple operations. It is also useful for discussing questions such as whether a given tropical linear space is an intersection of a given number of tropical hyperplanes, as the dual version of this questions is often easier to visualize.

Proposition 4.2.11. *The bounded part of $L^{\perp}$ is negative that of L. If L is a tropical linear space of dimension d in n space then the bounded part of L is at most* $\min(d, n-d)$*-dimensional.*

Proof. The first sentence follows because the property of being loop- and co-loop-free is self dual. L is d-dimensional and $L^{\perp}$ is $(n-d)$-dimensional; as the bounded part of L is isomorphic to a subcomplex of both it is at most $\min(d, n-d)$-dimensional. $\square$

4.3 Trop $G(3,6)$ and Matroidal Decompositions

Now that we have the terminology of matroidal subdivisions available to us, it is worthwhile to see how it can be used to described the various faces of Trop $G(3,6)$ in Section 3.3. There are three types of rays of Trop $G(3,6)$, which we referred to as types E, F and G. Each of these correspond to a minimal decomposition of $\Delta(3,6)$ into matroidal polytopes.

The ray E_{ijk} corresponds to a subdivision of $\Delta(3,6)$ into two polytopes: the first is the convex hull of the vertex e_{ijk} and all of its neighbors (as a polytope, this is isomorphic to the cone on the product of two triangles) and the second is the convex hull of all vertices except for e_{ijk}. As matroids, these two polytopes correspond to the point configurations in Figure 4.2.

The ray F_{ijkl} also corresponds to the splitting of $\Delta(3,6)$ into two polytopes. One is the convex hull of all of the vertices e_{abc} where $|\{a,b,c\} \cap \{i,j,k,l\}|$ is either 2 or 3, the other is the convex hull of all vertices e_{abc} where this intersection has size 1 or 2. These correspond to the dual matroid in Figure 4.3.

Finally, the ray G_{ijklmn} corresponds to the decomposition of $\Delta(3,6)$ into

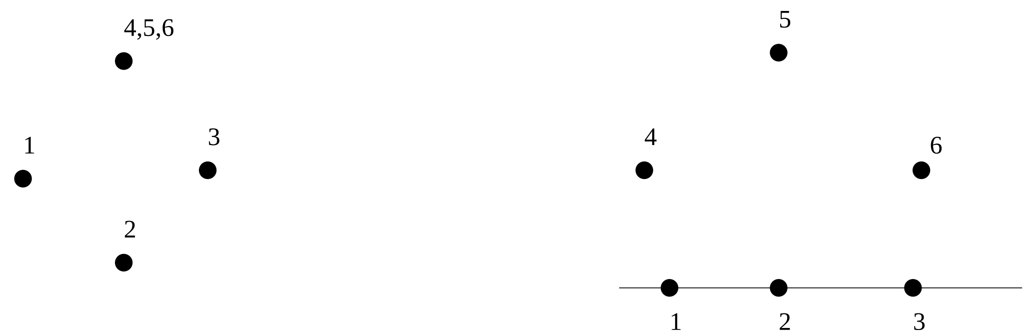

Figure 4.2: The two facets of the E-decomposition

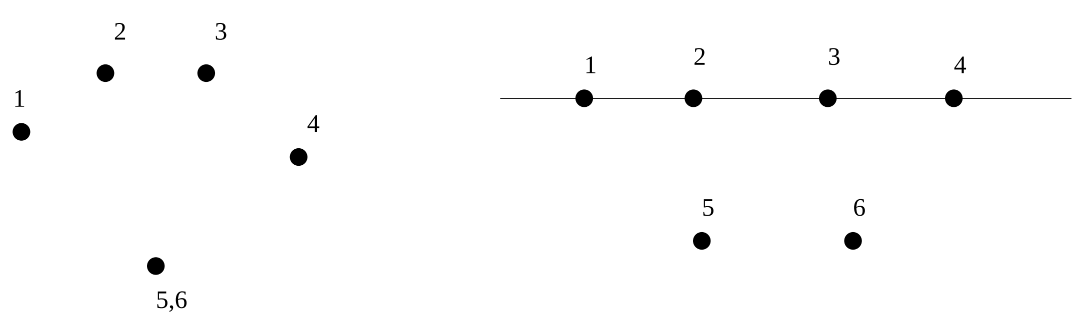

Figure 4.3: The two facets of the F-decomposition

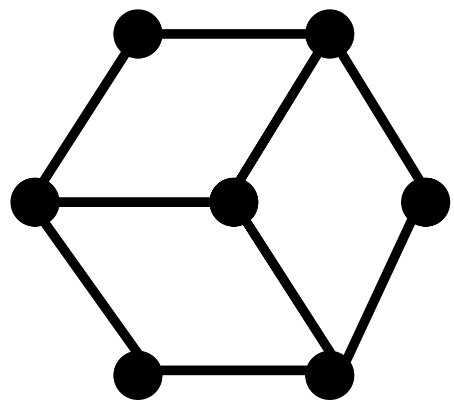

Figure 4.4: The hexagon subdivided into three rhombi

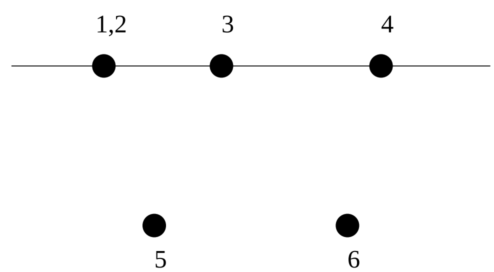

Figure 4.5: The facet (which occurs three times) of the G-decomposition

three polytopes. We will describe the case of g_{123456}, leaving the reader to use the S_6 symmetry of the problem to understand the other 29 cases. Consider the linear map ϕ which takes $\Delta(3,6)$ to the plane $x+y+z=3$ *via* $(a,b,c,d,e,f) \mapsto (a+b, c+d, e+f)$. The image of ϕ is a hexagon. The decomposition of $\Delta(3,6)$ corresponds to taking the preimage under ϕ of the decomposition of the hexagon into rhombi seen in Figure 4.4. Each of these rhombi corresponds to the matroid shown in Figure 4.5.

The various facets of Trop $G(3,6)$ then correspond to compatible ways of superimposing these decompositions. The type $FFFGG$, for example, can be described using the map ϕ of the previous paragraph as the decomposition of the hexagon into six equilateral triangles. Each of the vertices of the bipyramid corresponds to a different coarsening of this decomposition; see Figure 4.6. The reader may enjoy labelling the other faces of the bipyramid with intermediate decompositions of the hexagon.

It is possible to largely visualize the other facets of Trop $G(3,6)$ as pull-backs of decompositions of the hexagon as well. In each picture of Figure 4.7, we show a hexagon broken up into triangles and rhombi and decorated with some circular arcs. The meaning in each case is to pull back the decompositions into triangles and rhombi along ϕ, thus obtaining a decomposition of $\Delta(3,6)$. Then, for each circled vertex of the hexagon, we take a vertex e of $\Delta(3,6)$ that maps to this circled vertex under ϕ. We then slice off e and all of its neighbors as a seperate piece of the decomposition. In most cases, all choices of e are equivalent up to S_6 symmetry. The cases $EEFF1$ and $EEFF2$ are distinguished by the fact that, in type $EEFF1$, the two circled vertices are lifted to antipodal points of $\Delta(3,6)$ and in type $EEFF2$ they are lifted to non-antipodal points. In type $EEEE$, the four lifts must be chosen in the unique way (up to S_6 symmetry) that none of them are adjacent. The best way to describe this is that the remaining central region after these 4 vertices have been sliced off must correspond to the graphical matroid of K_4.

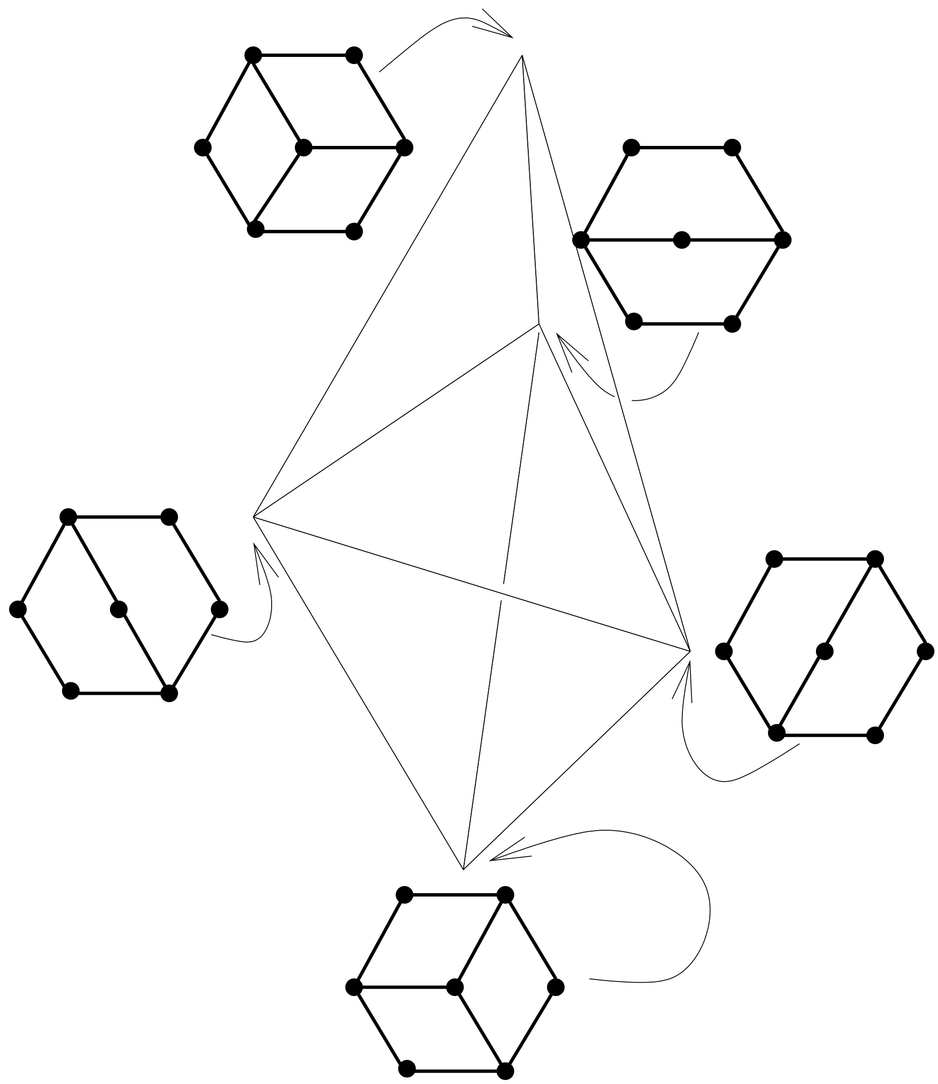

Figure 4.6: The faces of the bipyramid correspond to subdivisions of the hexagon

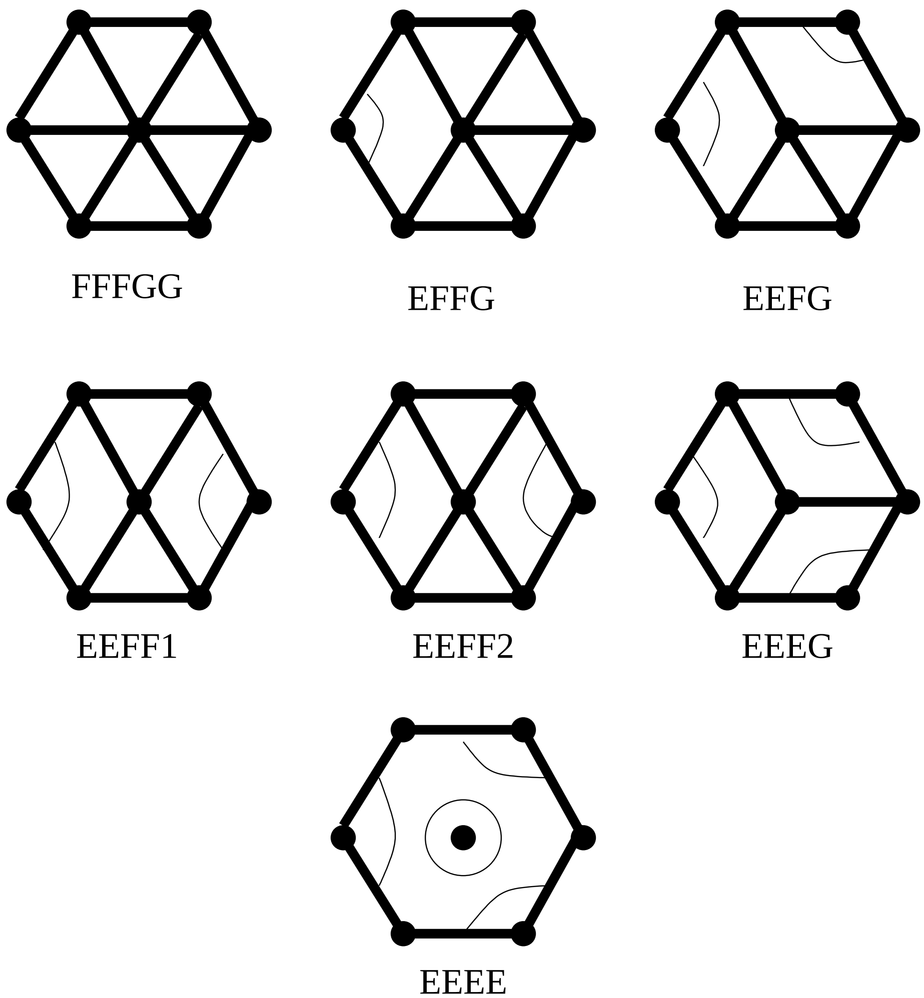

Figure 4.7: The facets of Trop $G(3,6)$

4.4 Stable Intersections

If L and L' are two ordinary linear spaces of dimensions d and d' in n space which meet transversely then $L \cap L'$ is a $(d + d' - n)$-dimensional space whose Plücker coordinates are given by

$$P_J(L \cap L') = \sum_{\substack{I \cap I' = J \\ |I| = d \\ |I'| = d'}} \pm P_I(L) P_{I'}(L').$$

Note that the summation conditions on I and I' guarantee that $I \cup I' = [n]$.

We will discuss a tropical version of this formula. If L and L' are tropical linear spaces of dimension d and d', define a real valued function $p(L \underset{\text{stable}}{\cap} L')$ on $\binom{[n]}{d+d'-n}$ by

$$p_J(L \underset{\text{stable}}{\cap} L') = \min_{\substack{I \cap I' = J \\ |I| = d \\ |I'| = d'}} (p_I(L) + p_{I'}(L')).$$

At the moment, $p_J(L \underset{\text{stable}}{\cap} L')$ is just defined as a formal symbol, there is not yet any tropical linear space called $L \underset{\text{stable}}{\cap} L'$. We now prove that there is such a space.

Proposition 4.4.1. $p(L \underset{\textit{stable}}{\cap} L')$ *is a tropical Plücker vector.*

Proof. Set P to be the polyhedron of points above the lower convex hull of the points $(p_{i_1 \dots i_d}(L), e_{i_1} + \dots + e_{i_d}) \in \mathbb{R} \times \Delta(d, n)$ and define $P' \subset \mathbb{R} \times \Delta(d', n)$ similarly. Include $\Delta(d + d' - n, n)$ into the Minkowski sum $\Delta(d, n) + \Delta(d', n)$ by $\iota : e \mapsto e + (1, \dots, 1)$. Set $Q = (P + P') \cap (\mathbb{R} \times \iota(\Delta(d + d' - n, n)))$. Let

$J \in \binom{[n]}{d+d'-n}$. The height of Q above the vertex $\iota(e_J)$ of $\iota(\Delta(d+d'-n,n)$ is

$$\min_{\substack{e_I \cap e_{I'} = \iota(e_J) \\ |I|=d \\ |I'|=d'}} p_I + p_{I'} = \min_{\substack{I \cap I' = J \\ |I|=d \\ |I'|=d'}} p_I + p_{I'}.$$

So we see that projecting the bottom faces of Q back down to $\Delta(d+d'-n,n)$ yields the regular subdivision of $\Delta(d+d'-n,n)$ induced by $p(L \underset{\text{stable}}{\cap} L')$.

So we see that every face of $\mathcal{D}_{p(L \underset{\text{stable}}{\cap} L')}$ is of the form

$$R := \iota(\Delta(d+d'-n,n)) \cap (P_M + P_{M'}) .$$

The vertices of this face are the points $\iota(e_J) \in \iota(\Delta(d+d'-n,n))$ where $J = I \cap I'$ for I and I' bases of M and M' with $I \cup I' = [n]$.

There is an operation $M \wedge M'$, called matroid intersection ([49], section 7.6), which takes two matroids of ranks d and d' and produces a third matroid. The spanning sets of $M \wedge M'$ are precisely the sets of the form $U \cap U'$ where U and U' span M and M' respectively. $M \wedge M'$ always has rank at least $d+d'-n$. If the rank of $M \wedge M'$ is larger than $d+d'-n$ then there are no pairs of bases I and I' of M and M' with $I \cup I' = [n]$. In this case R is empty and does not contribute a face to $\mathcal{D}_{p(L \underset{\text{stable}}{\cap} L')}$. Alternatively, the rank of $M \wedge M'$ is $d+d'-n$. Then we see that $R = P_{M \wedge M'}$. □

We have the immediate corollary

Corollary 4.4.2. *Every face of* $\mathcal{D}_{p(L \underset{\textit{stable}}{\cap} L')}$ *is of the form* $P_{M \wedge M'}$ *for* P_M *and* $P_{M'}$ *faces of* $\mathcal{D}_{p(L)}$ *and* $\mathcal{D}_{p(L')}$.

We thus see that there is a well defined tropical linear space $L \underset{\text{stable}}{\cap} L'$. Our aim now is to justify the notation $L \underset{\text{stable}}{\cap} L'$ by showing that $L \underset{\text{stable}}{\cap} L'$ is the "stable" intersection of L and L'. We first need a definition and a combinatorial lemma:

Let M and M' be two loop-free matroids on $[n]$ with connected decomposition $M = \bigoplus_{i \in I} M_i$ and $M' = \bigoplus_{i \in I'} M'_i$. Let S_i and S'_i be the ground sets of M_i and M'_i. Let $\Gamma(M, M')$ be the bipartite graph (possibly with multiple edges) whose vertex set is $I \sqcup I'$ and which has an edge between i and i' for each element of $S_i \cap S'_{i'}$. We say that M and M' are *transverse* if $\Gamma(M, M')$ is a forest without multiple edges.

Lemma 4.4.3. *Suppose that M and M' are transverse and let $\Gamma = \Gamma(M, M')$. Then $M \wedge M'$ is a loop-free matroid whose connected components are in bijection with the components of Γ. $M \wedge M'$ is the matroid formed by repeatedly taking parallel connections along the edges of Γ. $M \wedge M'$ has rank $d + d' - n$.*

The phrase "repeatedly taking parallel connections" deserves a proper definition, which will unfortunately be somewhat lengthy. What we mean is the following: let F be a forest, let E be a finite set and let each edge of F be labeled with an element of E, no two edges receiving the same label. Suppose that, for each vertex $v \in F$, we are given a subset $S(v)$ of E. We require that $S(v) \cap S(w) = \{e\}$ if e is the edge (v, w), that $S(v) \cap S(w) = \emptyset$ if there is no edge joining v and w and that $\bigcup_{v \in F} S(v) = E$. For every $v \in F$, we place the structure of a connected loop-free matroid $M(v)$ on $S(v)$. We then define a matroid N on the ground set E by the following procedure:

1. If F has no edges, let $N = \bigoplus_{v \in F} M(v)$.

2. Otherwise, let v and w be vertices of F joined by an edge e.

3. Form a new forest F' where e is contracted to a single vertex u. Set $S'(u) = S(v) \cup S(w)$ and let $M'(u)$ be the parallel connection of $M(v)$ and $M(w)$ along e. For all $t \in F$ other than u, set $S'(t) = S(t)$ and $M'(t) = M(t)$.

4. Return to step (1) with the new forest F' and the new $S'(\cdot)$ and $M'(\cdot)$.

The claim is that (a) N does not depend on the order of the contractions and (b) $N = M \wedge M'$ when we take $F = \Gamma$, $S(i) = S_i^{(')}$ and $M(i) = M_i^{(')}$. Here $S_i^{(')}$ is shorthand for "S_i if $i \in I$ and S'_i if $i \in I'$ and similarly for $M_i^{(')}$.

Proof. Let T, $S(\cdot)$ and $M(\cdot)$ be as above and let $N(F, S(\cdot), M(\cdot))$ be the matroid constructed by the above procedure. We claim that $U \subset E$ is a basis of $N(F, S(\cdot), M(\cdot))$ if and only if there exist bases $B(v)$ of $M(v)$ such that, for every $e \in E$, precisely one of the following four conditions holds:

1. $e \in U$, $e = (v, w)$ is an edge of F and $e \in B(v) \cap B(w)$.

2. $e \in U$, $e \in S(v)$ is not an edge of F and $e \in B(v)$.

3. $e \notin U$, $e = (v, w)$ is an edge of F and e lies in precisely one of $B(v)$ and $B(w)$.

4. $e \notin U$, $e \in S(v)$ is not an edge of F and $e \notin B(v)$

This is the definition of direct sum when F has no edges and it is easy to check that this description is unchanged by contracting an edge, so this is indeed a correct description of the bases of $N(F, S(\cdot), M(\cdot))$. Note that this description is independent of the order of contractions, so we see that $N(T, S(\cdot), M(\cdot))$ is independent of the order of contraction.

Applying the above description in the case where $F = \Gamma(M, M')$ *etc.*, we see that the bases of $N(\Gamma, M(\cdot), S(\cdot))$ are precisely the sets of the form $B \cap B'$ where B and B' are bases of M and M' with $B \cup B' = [n]$. If M and M' are such that there exist bases B and B' with $B \cup B' = [n]$ then the bases of $M \wedge M'$ are exactly the sets of the form $B \cap B'$ for such (B, B'). But the iterative construction above clearly does produce a matroid and every matroid has at least one basis, so such B and B' do exist. $|B \cap B'| = |B| \cup |B'| - |B \cup B'|$ so N has rank $d + d' - n$. Finally, observe that parallel connections preserves loop-freeness, so N is loop-free. $\square$

Proposition 4.4.4. *Let $0 \leq d, d' \leq n$ with $d + d' \geq n$. Suppose that L and L' are tropical linear spaces in n-space of dimension d and d'. Then $L \underset{stable}{\cap} L' \subseteq L \cap L'$. If L and L' meet transversely then $L \underset{stable}{\cap} L' = L \cap L'$.*

Proof. Let $\mathcal{D}$, $\mathcal{D}'$ and $\mathcal{E}$ be the subdivisions of $\Delta(d,n)$, $\Delta(d',n)$ and $\Delta(d+d'-n,n)$ corresponding to L, L' and $L \underset{\text{stable}}{\cap} L'$. Let $\mathcal{D}^\vee$, $(\mathcal{D}')^\vee$ and $\mathcal{E}^\vee$ denote the dual subdivisions of $\mathbb{R}^n$.

Consider a particular $w \in \mathbb{R}^n$, we need to show that $w \in L \underset{\text{stable}}{\cap} L'$ implies that $w \in L$ and $w \in L'$ and that the reverse holds if L and L' meet transversely. Let $M^\vee$, $(M')^\vee$ and $N^\vee$ be the faces of $\mathcal{D}^\vee$, $(\mathcal{D}')^\vee$ and $\mathcal{E}^\vee$ respectively containing w and let P_M, $P_{M'}$ and P_N be the respective dual faces of $\mathcal{D}$, $\mathcal{D}'$ and $\mathcal{E}$. Then $P_N = (P_M + P_{M'}) \cap \iota(\Delta(d+d'-n,n))$. We must show that, if M or M' has a loop then N has a loop and that, if L and L' meet transversely and N has a loop then M or M' does as well.

First, suppose that M contains a loop e. Then, for every $x \in P_M$, we have $x_e = 0$. For every $x \in P_{M'}$ we have $x_e \leq 1$. Thus, for $x \in P_M + P_{M'}$, $x_e \leq 1$ and thus $x_e = 0$ on P_N. So e is a loop of N.

Now assume that L and L' meet transversely. Let $M = \bigoplus_{i \in I} M_i$ and $M' = \bigoplus_{i \in I'} M'_i$ be the components of M and M' and let S_i and S'_i be the ground sets of M_i and M'_i respectively.

Lemma 4.4.5. *M and M' are transverse,* i.e., *$\Gamma(M,M')$ is a forest and has no multiple edges.*

Proof. Let $|I| = e$ and $|I'| = e'$, so e and e' are the dimensions of $M^\vee$ and $(M')^\vee$. The affine linear space spanned by $M^\vee$ is cut out by the equations $x_i - x_j =$ constant when edges i and j have the same endpoint in I, and similarly for

$(M')^\vee$. Thus, $M^\vee \cap (M')^\vee$ spans the affine linear space cut out by the equations $x_i - x_j =$ constant whenever edges i and j are in the same component of Γ. We assumed that L and L' meet transversely, so the dimension of the last linear space must be $e + e' - n$ and Γ must have $e + e' - n$ connected components. A graph with $e + e'$ vertices, n edges and $e + e' - n$ connected components must be a forest without multiple edges. $\square$

We now know by Lemma 4.4.3 that $M \wedge M'$ is loop-free. $\square$

The next theorem explains the motivation for the notation $L \underset{\text{stable}}{\cap} L'$ and the geometric meaning of $L \underset{\text{stable}}{\cap} L'$ when L and L' are not transverse.

Theorem 4.4.6. *Let L and L' be tropical linear spaces. Then, for a generic $v \in \mathbb{R}^n$, the linear spaces L and $L' + v$ meet transversely and $L \underset{\textit{stable}}{\cap} L' = \lim_{v \to 0}(L \cap (L' + v))$ where the limit is taken through generic v.*

Proof. Let F and F' be faces of L and L'. Suppose that the affine linear spaces spanned by F and F' together fail to span $\mathbb{R}^n$. Then, for v outside a proper subspace of $\mathbb{R}^n$ we have $F \cap (F' + v) = \emptyset$. Thus, for a generically chosen v, every such F and $F' + v$ fail to meet. The only faces of L and $L' + v$ that do meet, then, meet transversely.

We have $p_I(L' + v) = p_I(L') + \sum_{i \in I} v_i$ so the tropical Plücker coordinates of $L' + v$ vary continuously with v. The tropical Plücker coordinates of $L \underset{\text{stable}}{\cap} (L' + v)$ similarly vary continuously with v. When v is chosen generically, L and

$L' + v$ meet transversely so $L \cap (L' + v) = L \underset{\text{stable}}{\cap} (L' + v)$. Take limits of both sides as $v \to 0$ and use the proceeding continuity arguments to conclude that $\lim_{v \to 0}(L \cap (L' + v)) = L \underset{\text{stable}}{\cap} L'$. $\square$

4.5 Realizability of Tropical Linear Spaces

In this section, we will discuss the question of which tropical linear spaces are tropicalizations of actual linear spaces. In this section, we will denote actual linear spaces (over K or κ) by boldface characters such as $\boldsymbol{L}$.

Let $\boldsymbol{L}$ be a linear space with Plücker coordinates $P_{i_1 \ldots i_d} \neq 0$. Let $p_{i_1 \ldots i_d} = v(P_{i_1 \ldots i_d})$. The P's obey the Plücker relations, so the p's obey the tropical Plücker relations. Our next proposition shows that the combinatorially defined $L(p)$ truly does reflect the geometrically defined $\operatorname{Trop} \boldsymbol{L}$.

Proposition 4.5.1. *We have* $\operatorname{Trop} \boldsymbol{L} = L(p)$.

Proof. Every point $(x_1, \ldots, x_n)$ in $\boldsymbol{L}$ obeys $\sum (-1)^r P_{j_1 \ldots \widehat{j_r} \ldots j_{d+1}} X_{j_r} = 0$ so the $v(X_i)$ lie in $\operatorname{Trop}(\sum_{r=1}^{d+1} (-1)^r P_{j_1 \cdots \hat{j_r} \cdots j_{d+1}} X_{j_r})$. Thus, $\operatorname{Trop} \boldsymbol{L} \subseteq L(p)$.

For the converse direction, suppose $w \in L(p)$, we will prove that $w \in \operatorname{Trop} \boldsymbol{L}$. Without loss of generality, let $w = 0$ and suppose that $0 = \min_{I \in \binom{[n]}{d}} p_I$. Let $P_I(0) \in \kappa$ be the image of P_I in $\kappa = R/\mathcal{M}$. Clearly, the $P_I(0)$ obey the Plücker relations and hence correspond to an ideal $\boldsymbol{L(0)} \subset \kappa^n$. It is easy to see that $\boldsymbol{L(0)} = \operatorname{in}_w \boldsymbol{L}$. Let M be the set of I for which $p_I = 0$. By assumption, M is the set of bases of a loop free matroid. So $\boldsymbol{L(0)}$ is not contained in any of the coordinate planes of κ^n and thus the corresponding ideal contains no monomial. $\square$

We can now prove Theorem 4.1.1.

Theorem 4.5.2. *The points of $\mathcal{G}'_{d,n}$ parameterize the possible subsets of $\mathbb{R}^n$ that can occur as $\operatorname{Trop} \boldsymbol{L}$ for $\boldsymbol{L}$ a d-dimensional linear space in K^n with all Plücker coordinates nonzero.*

Proof. First, suppose that L is such a linear space with Plücker coordintes $P_I \neq 0$. Then $\operatorname{Trop} \boldsymbol{L} = L(v(P))$, by the above proposition. By Proposition 4.2.7, $L(v(P))$ determines $v(P)$ up to an additive constant. Thus, we see that $\operatorname{Trop} \boldsymbol{L}$ determines a point of $\operatorname{Trop} G(d,n)/(1,\ldots,1) = \mathcal{G}'_{d,n}$.

On the other hand, if p is a point of $\operatorname{Trop} G(d,n)$ then $p_I = v(P_I)$ for P_I the actual Plücker coordinates of some linear space $\boldsymbol{L}$. Then $L(p)$ is determined by p and the above proposition shows that $\operatorname{Trop} \boldsymbol{L} = L(p)$. So every point $\mathcal{G}_{d.n}$ does give rise to a polyhedral complex in $\mathbb{R}^n$ of the form $\operatorname{Trop} \boldsymbol{L}$. The two maps are easily seen to be inverse. $\square$

It also turns out that $\underset{\text{stable}}{\cap}$ does actually capture the behavior of linear spaces under intersection.

Proposition 4.5.3. *Let $\boldsymbol{L}$ and $\boldsymbol{L'} \subset K^n$ be two linear spaces. Then there exists a linear space $\tilde{\boldsymbol{L}}' \subset K^n$ with $\operatorname{Trop} \tilde{\boldsymbol{L}}' = \operatorname{Trop} \boldsymbol{L'}$ and $\operatorname{Trop} \boldsymbol{L} \underset{stable}{\cap} \operatorname{Trop} \boldsymbol{L'} = \operatorname{Trop} \boldsymbol{L} \underset{stable}{\cap} \operatorname{Trop} \tilde{\boldsymbol{L}}' = \operatorname{Trop}(\boldsymbol{L} \cap \tilde{\boldsymbol{L}}')$. If $\operatorname{Trop} \boldsymbol{L}$ and $\operatorname{Trop} \boldsymbol{L'}$ meet transversely,* *then we already have* $\operatorname{Trop} \boldsymbol{L} \underset{stable}{\cap} \operatorname{Trop} \boldsymbol{L'} = \operatorname{Trop}(\boldsymbol{L} \cap \boldsymbol{L'})$ *without having to choose an* $\tilde{\boldsymbol{L}}'$*.*

Proof. Let d and d' be the dimensions of $\boldsymbol{L}$ and $\boldsymbol{L'}$. Let P_I, P'_I and Q_I be the Plücker coordinates of $\boldsymbol{L}$, $\boldsymbol{L'}$ and $\boldsymbol{L} \cap \boldsymbol{L'}$, so

$$Q_J = \sum_{\substack{I \cap I' = J \\ |I| = d \\ |I'| = d'}} \pm P_I P'_{I'}.$$

If q_I are the tropical Plücker coordinates of $\operatorname{Trop} \boldsymbol{L} \underset{\text{stable}}{\cap} \operatorname{Trop} \boldsymbol{L'}$, we have

$$q_J = \min_{\substack{I \cap I' = J \\ |I| = d \\ |I'| = d'}} v(P_I P'_{I'}).$$

We want to have $q_J = v(Q_J)$, so we will be successful if there is no cancellation of leading terms in $\sum \pm P_I P'_{I'}$.

Write $U = \{x \in K^* : v(x) = 0\}$ and let $(u_1, \ldots, u_n) \in U^n$. Let $\tilde{\boldsymbol{L}}' = \operatorname{diag}(u_1, \ldots, u_n)\boldsymbol{L'}$. Letting $\tilde{P'}_I$ denote the Plücker coordinates of $\tilde{\boldsymbol{L}}'$, we have $\tilde{P'}_I = P'_I \prod_{i \in I} u_i$. When the u_i are chosen generically, there is no cancellation of leading terms in the sum $\sum \pm P_I \tilde{P'}_{I'} = \sum \pm P_I P'_I \prod_{i \in I'} u_i$. For a generic u, therefore, $\tilde{\boldsymbol{L}}'$ has the desired property.

We now must prove the second claim, that if $\operatorname{Trop} \boldsymbol{L}$ and $\operatorname{Trop} \boldsymbol{L'}$ meet transversely then we never have cancellation in this sum. We will do this by showing that one term has a lower valuation than all of the others. Let $\mathcal{E}$ be the subdivision of $\Delta(d + d' - n, n)$ induced by q and let $\mathcal{D}$ and $\mathcal{D}'$ be the subdivisions of $\Delta(d, n)$ and $\Delta(d', n)$ corresponding to $\operatorname{Trop} \boldsymbol{L}$ and $\operatorname{Trop} \boldsymbol{L'}$. Let P_N be a facet of $\mathcal{E}$ containing the vertex J and let $w \in N^\vee$. Let $M^\vee$ and $(M')^\vee$ be the faces of $\mathcal{D}^\vee$ and $(\mathcal{D}')^\vee$ containing w. Without loss of generality, we can take $w = 0$ and $0 = \min p_I = \min p'_I$.

Let $\bigoplus_{i \in I} M_i$ and $\bigoplus_{i \in I'} M'_i$ be the connected components of M and M'. Write $\Gamma = \Gamma(M, M')$. As in Lemma 4.4.5, we see that Γ is a forest. The fact that J is a basis of N indicates that there are B and B' with $B \cap B' = J$ and $B \cup B' = [n]$ such that, for each M_i, the set $B \cap M_i$ a basis of M_i and similarly for B'. If we show there is only one such pair (B, B'), that will indicate that there is only one term in the sum with valuation zero and hence no cancellation.

Suppose there were two such pairs, (B_1, B'_1) and (B_2, B'_2). For $e \in [n]$ and $S \subset [n]$, define $[e \in S]$ to be 1 if $e \in S$ and 0 otherwise. Note that $[e \in B_s] + [e \in B'_s] = 1 + [e \in J]$. Set

$$\begin{aligned} \alpha(e) &= [e \in B_1] - [e \in B'_1] - [e \in B_2] + [e \in B'_2] \\ &= 2\left([e \in B_1] - [e \in B_2]\right). \end{aligned}$$

We think of α as a cochain in $C^1(\Gamma)$. We have

$$\begin{aligned} \partial\alpha(M_i) &= 2 \sum_{e \in M_i} \left([e \in B_1] - [e \in B_2]\right) \\ &= 2(|B_1 \cap M_i| - |B_2 \cap M_i|). \end{aligned}$$

Since $B_1 \cap M_i$ and $B_2 \cap M_i$ are both bases of M_i they have the same cardinality and $\partial\alpha(M_i) = 0$ for every M_i. Similarly, $\partial\alpha(M'_i) = 0$. So α is a cochain. But Γ is acyclic, so $\alpha = 0$. This in turn implies that $[e \in B_1] = [e \in B_2]$ for every

$e \in [n]$ so $B_1 = B_2$ and similarly $B'_1 = B'_2$. □

Having seen that every linear space over K gives rise to a tropical linear space, one might ask whether every tropical linear space arises in this manner. The answer to this question is a dramatic "no" and we conclude this section by showing various manners in which this can fail. We will say that a tropical linear space is *realizable* if it arises from a linear space over K.

Example 4.5.4. First, $\mathcal{D}$ can contain polytopes corresponding to non-realizable matroids; such a polytope obviously can not occur in a realizable decomposition. For example, if M is the non-Pappus matroid, removing P_M from $\Delta(3,9)$ leaves 8 matroidal polytopes, each of them a cone on $\Delta_2 \times \Delta_5$. Cutting $\Delta(3,9)$ into the non-Pappus matroid and the eight other pieces gives a non-realizable matroidal subdivision. More generally, any matroid at all may appear as a piece of a matroidal subdivision:

Proposition 4.5.5. *If M is a rank d matroid on n elements, the function $-\rho(\cdot)$ on $\binom{[n]}{d}$ obeys the tropical Plücker relations, where $\rho(S)$ is the rank of the flat spanned by S for any $S \subset [n]$. M is a face of the corresponding subdivision.*

Proof. Let $S \in \binom{[n]}{d-2}$ and let i, j, k and l be distinct elements of $[n] \setminus S$. Write $\rho_{M/S}$ for the rank function of M/S, we have $\rho(Sij) = \rho_{M/S}(ij) + \rho(S)$. Thus, we only need to check that $-\rho$ obeys the tropical Plücker relations in the case where $d = 2$ and $n = 4$, which is straight forward. M is a face of the corresponding subdivision because $-\rho$ is minimal on the vertices of M. □

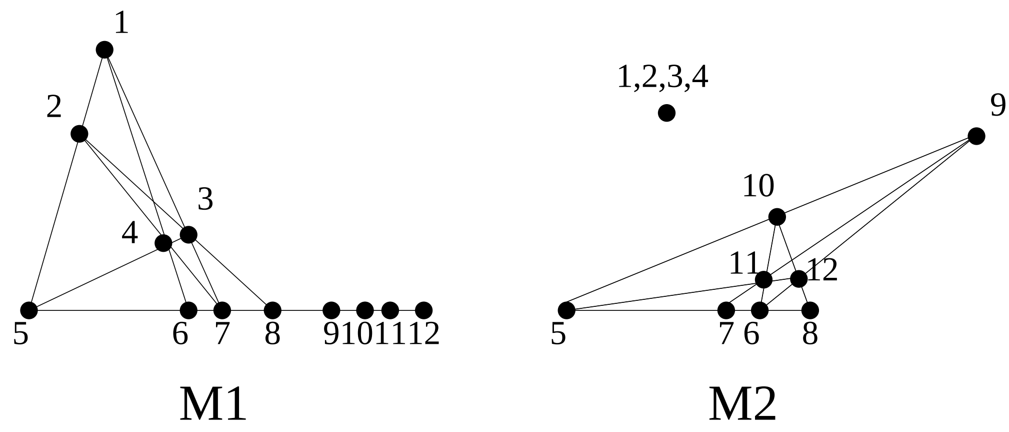

Figure 4.8: The matroids M_1 and M_2

Example 4.5.6. Even if every face of $\mathcal{D}$ is realizable, it is still possible for the subdivision to be non-realizable for global reasons. Suppose κ does not have characteristic 2 or 3. Let M_1 and M_2 be the rank 3 matroids on 12 points corresponding to the plane geometries in Figure 4.8. After removing M_1 and M_2 from $\Delta(3, 12)$, the remainder can be cut into cones on $\Delta_2 \times \Delta_8$. This gives a matroidal subdivision into realizable matroids. This subdivision is not realizable; if $p_{ijk} \in K$ where the Plücker coordinates of a linear space over K realizing this subdivision then the presence of M_1 implies that $p_{156}p_{178} = p_{167}p_{158}(1 +$ higher order terms). The presence of M_2 implies that $p_{156}p_{178} = p_{167}p_{158}(-2 +$ higher order terms), a contradiction.

Example 4.5.7. It is possible to have a matroidal subdivision that is not regular and hence certainly can not be realizable. Map $\Delta(6, 12)$ to $\mathbb{Z}^3/(1, 1, 1)$ by $(x_1, \ldots, x_{12}) \to (x_1 + \ldots + x_4, x_5 + \ldots + x_8, x_9 + \ldots + x_{12})$. Subdivide $\Delta(6, 12)$

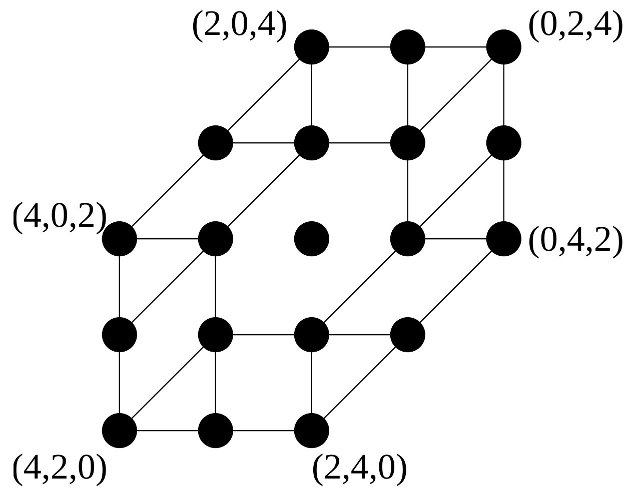

Figure 4.9: A non-regular subdivsion that induces a non-regular subdivision of $\Delta(6, 12)$.

according to the preimage of Figure 4.9. The resulting subdivision is not regular because Figure 4.9 is not.

Example 4.5.8. It is possible for a matroidal subdivision to be induced by two different p's such that one of the p's is realizable and the other is not. We give an example when K has characteristic 2, but there are characteristic zero examples. This will also serve as a good opportunity to demonstrate that $\mathcal{G}_{3,7}$ depends on the characteristic of κ.

Let F denote the Fano plane (see Figure 4.10), and let $\mathcal{D}$ be the subdivision of $\Delta(3, 7)$ into P_F and 7 cones on $\Delta_2 \times \Delta_3$. Let $C \subset \binom{[7]}{3}$ be the set of three-element circuits of F.

Proposition 4.5.9. *If* $p : \binom{[7]}{3} \to \mathbb{R}$ *induces the subdivision* $\mathcal{D}$ *then* p *is a tropical Plücker vector. If* κ *does not have characteristic* 2 *then* $p \notin \mathcal{G}_{3,7}$*. If*

K has characteristic 2 then p may or may not be in $\mathcal{G}_{3,7}$. More specifically, we can use the translation by $\phi(\mathbb{R}^7)$ to assume that $p_I = 0$ for $I \notin C$. With this normalization, p induces $\mathcal{D}$ if and only if $p_I > 0$ for $I \in C$ and $p \in \mathcal{G}_{3,7}$ if and only if the minimum of p_I, $I \in C$, is not unique.

Remark: Recall that we assumed that K and κ have the same characteristic, so this proposition does cover all cases. We have deliberately phrased it in order not to refer to the possibility that K has characteristic 0 and κ has characteristic 2. In this setting, $\mathcal{G}_{3,7}$ would be a polyhedral complex, not a fan. Using the normalization in the proposition that $p_I = 0$ for $I \notin C$, the portion of $\mathcal{G}_{3,7}$ which induces $\mathcal{D}$ is given by the equation that the minimum value of p_I for $I \in C$ should be not unique and less than $v(2)$.

Proof. Certainly, if the characteristic of κ is not 2 then $p \notin \mathcal{G}_{3,7}$, as the Fano plane is not realizable over κ. Also, p is a tropical Plücker vector, as the faces of $\mathcal{D}$ are matroidal.

In characteristic 2, the following relation holds among the Plücker coordinates:

$$\begin{aligned} 0 \quad = \quad & P_{357}P_{457}\underline{P_{126}} + P_{157}P_{567}\underline{P_{234}} + P_{125}P_{457}\underline{P_{367}} \\ & + P_{127}P_{357}\underline{P_{456}} + P_{235}P_{567}\underline{P_{147}} + P_{136}P_{457}\underline{P_{257}} + P_{267}P_{475}\underline{P_{135}} + P_{137}\underline{P_{257}P_{456}}. \end{aligned}$$

Here the underlined variables are those corresponding to the elements of C. We see that, with our assumption that p is zero on all vertices outside C, the valua-

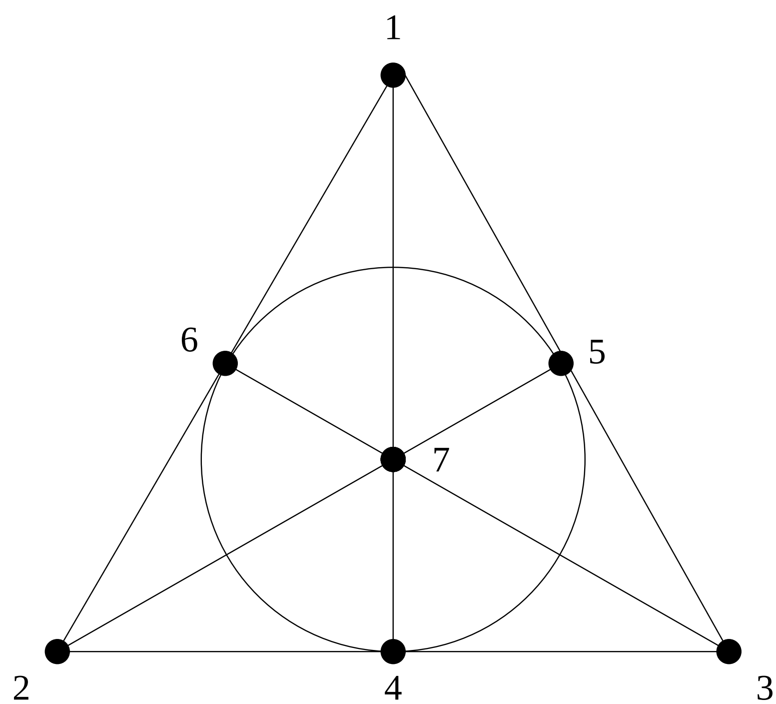

Figure 4.10: The Fano Plane

tions of the various terms are p_I as I runs over C and also $p_{257}+p_{456}$. As p_{257} and $p_{456}>0$, this last term can't be minimal. We see that the p's must obey the relation that the minimum of p_I for $I \in C$ is not unique in order to be realizable.

Finally, we must show that the condition that $\min_{I\in C} p_I$ not be unique is also sufficient. By the symmetry of the problem, we may assume that $\min_{I\in C} p_I$ is attained at p_{456} as well as at P_I for some $I \in C \setminus \{456\}$. Consider

$$\boldsymbol{L} = \operatorname{RowSpan}\begin{pmatrix} 1 & 0 & 0 & at^{p_{234}} & 1 & 1+ft^{p_{367}} & 1 \\ 0 & 1 & 0 & 1+dt^{p_{147}} & bt^{p_{135}} & 1 & 1 \\ 0 & 0 & 1 & 1 & 1+et^{p_{257}} & ct^{p_{126}} & 1 \end{pmatrix}.$$

Here a, b, c, d, e and f are generic members of $U=\{x \in K^* : v(x)=0\}$.

For $I \notin C$, we have $v(P_I(\boldsymbol{L}))=0$. For $I \in C \setminus \{(456)\}$, we have $p_I(\boldsymbol{L}) = ut^{p_I}$, for some $u \in U$, so $v(P_I(L)) = p_I(L)$. Finally,

$$v(P_{456}) = v\left(at^{p_{234}} + bt^{p_{135}} + ct^{p_{126}} + dt^{p_{147}} + et^{p_{257}} + ft^{p_{367}} + \cdots\right)$$

where the remaining eight terms hidden in the "$\cdots$" each have valuation at least as large as $2\min_{I\in C\setminus\{(456)\}} p_I$. If a, b, c, d, e, f are chosen generically, there is no cancellation, so we have $v(P_{456}) = \min_{I\in C\setminus\{(456)\}} p_I = p_{456}$ as desired. ☐

We end by discussing two naturally occurring examples of tropical Plücker coordinates for which I do not know whether or not the corresponding tropical linear space is realizable.

Problem 4.5.10. Let T be a tree with leaves labeled by $[n]$ and a positive weight on each edge. For $I \subset [n]$, define $[I]$ to be the minimal subtree of T containing

I and set p_I to be negative the sum of the weights of all edges in $[I]$. In [31], Lior Pachter and I observe that p_I are tropical Plücker coordinates; are they realizable?

Problem 4.5.11. Let $R = \mathbb{R}((t))$. Then R is an ordered field where $at^\alpha + \cdots$ is positive if and only if a is. Thus, it makes sense to define a $d \times d$ matrix A with entries in R to be positive definite if $vAv^T > 0$ for all $v \in R^d \setminus \{0\}$. Let $X_1, \ldots, X_n$ be $d \times d$ positive definite symmetric matrices with entries in R. Set $P_{i_1 \ldots i_d}$ equal to the coefficient of $x_{i_1} \cdots x_{i_d}$ in $\det \sum x_i X_i$. The P_I do not obey the Plücker relations. Nonetheless, I show in [37] that $v(P_I)$ do obey the tropical Plücker relations! Are the $v(P_i)$ realizable? If so, give a natural linear space over R, constructed from the X_i, that realizes them.

4.6 Series-Parallel Matroids and Linear Spaces

Recall that if G is a connected graph (finite, possibly with multiple edges or loops) then the associated *graphical matroid* is a matroid whose elements are the edges of G, whose bases are the spanning trees and whose circuits are the circuits. See [49], section 6.1 for an overview of graphical matroids. The rank of this matroid is the number of vertices of G minus one. A matroid is called *series-parallel* if it corresponds to a series-parallel graph, *i.e.* a graph which can made by starting from a single edge connecting two distinct vertices by repeatedly applying the series and parallel extension operators. (See Figure 4.11.) For a

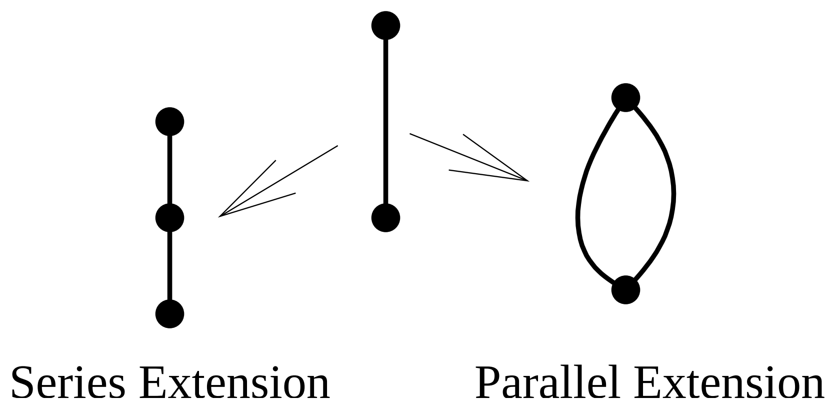

Figure 4.11: Series and Parallel Extension

general introduction to series-parallel matroids, see section 6.4 of [49].

We will call a matroidal decomposition $\mathcal{D}$ of $\Delta(s,n)$ a *series-parallel decomposition* if every facet of $\mathcal{D}$ is the polytope of a series-parallel matroid and the associated tropical linear space will be called a *series-parallel tropical linear space*. We will see that *series-parallel tropical linear spaces* are the most natural and most manageable tropical linear spaces.

Recall that one of our main aims is to prove:

The f-Vector Conjecture. Every d-dimensional tropical linear space L in n-space has at most $\binom{n-2i}{d-i}\binom{n-i-1}{i-1}$ bounded faces of dimension i, with equality if and only if L is series-parallel.

We spend the rest of this section presenting what appears to be a new way of treating series-parallel matroids. These methods will be applied in sections 4.8 and 4.9.

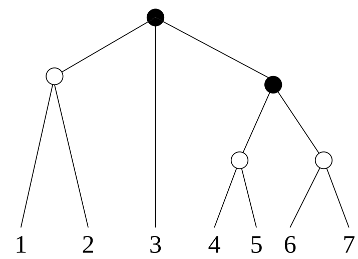

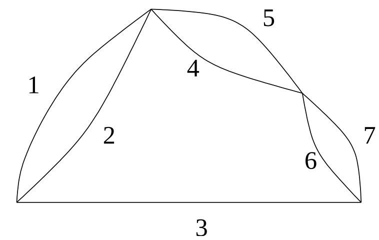

Figure 4.12: A bi-colored tree and the corresponding matroid

Let T be a trivalent tree (*i.e.* every internal vertex of T has three neighbors) with n leaves whose $n-2$ internal vertices are colored white and black with $d-1$ black vertices and $n-d-1$ white vertices. For every edge e of T let $A_e \subset [n]$ be the leaves on one side of e and let a_e be the number of black vertices on that same side of e. Let $\Pi(T)$ be the subpolytope of $\Delta(d,n)$ consisting of all (x_i) such that, for every edge e, $a_e \leq \sum_{i\in A_e} x_i \leq a_e + 1$.

Proposition 4.6.1. $\Pi(T) = P_{\mu(T)}$ *for a series-parallel matroid* $\mu(T)$*. Every series-parallel matroid can be written as* $\mu(T)$ *for some bi-colored trivalent tree* T*.*

When it is necessary to emphasize the dependence of $\Pi(T)$ or $\mu(T)$ on the coloring c, we write $\Pi(T,c)$ or $\mu(T,c)$. Figure 4.12 shows a bi-colored tree and the assosciated matroid. Note that $\mu(T,c)$ does not determine T and c. The precise statement is that $\mu(T,c)$ determines and is determined by the non-trivalent bicolored tree which is produced from T by contracting all edges between vertices of the same color.

Proof. Our proof is by induction on n; when $n = 2$ the result is clear. Let i and $j \in [n]$ be two leaves of T that border a common vertex v. Define a new colored trivalent tree T' by deleting i and j from T and forgetting the coloring of v, which is now a leaf. By induction, $\Pi(T') = P_{\mu(T')}$ for a series-parallel matroid $\mu(T')$.

Case I: v is colored white in T. Then, cutting T at the edge separating i and j from the rest of T, we see that $x_i + x_j \leq 1$. We define a map $\phi : \mathbb{R}^n \to \mathbb{R}^{n-1}$ by $\phi(x)_v = x_i + x_j$ and $\phi(x)_k = x_k$ otherwise. An easy inspection of the defining inequalities shows that $\Pi(T) = \Delta(d, n) \cap \phi^{-1}(\Pi(T'))$. This precisely says that $\Pi(T)$ is the polytope associated to a series extension of $\mu(T')$ by v.

Case II: v is colored black in T. This is just like the other case except that $x_i + x_j \geq 1$, we define $\phi(x)_v = x_i + x_j - 1$ and we get a parallel extension.

To show that every series-parallel matroid occurs in this way, reverse the argument. □

Not all of the inequalities above are necessary to define $\Pi(T)$. The next proposition identifies the facets of $\Pi(T)$.

Proposition 4.6.2. *The facets of* $\Pi(T)$ *are given by (1) the equations* $\sum_{i \in A_e} x_e \geq a_e$ *when* e *is an edge connecting two vertices of the opposite color and* A_e *is the set of leaves on the same side of* e *as the black endpoint, (2) the equations* $x_i \geq 0$ *when* i *is joined to a white vertex and (3) the equations* $x_i \leq 1$ *when* i *is joined to a black vertex. (All of these are subject to already having the equality* $\sum x_i = d$*.)*

Proof. We first show that all other equations are redundant. Let e be an internal edge with A_e and B_e the leaves on each side of e and a_e and b_e the number of black vertices on each side. Then $\sum_{i\in A_e} x_i = d - \sum_{i\in B_e} x_i$ so $\sum_{i\in A_e} x_i \leq a_e + 1$ is equivalent to $\sum_{i\in B_e} x_i \geq d-(a_e+1) = b_e$. If i is a leaf joined to a white vertex v, let e and e' be the other two edges issuing from v and let A_e and $A_{e'}$ be the sides of e and e' not containing v. Then

$$d = x_i + \sum_{j\in A_e} x_j + \sum_{j\in A_{e'}} x_j \geq x_i + a_e + a_{e'} = x_i + d - 1$$

implying that $x_i \leq 1$. Similarly, if i is a leaf connected to a black vertex, then the other inequalities imply $x_i \geq 0$.

Suppose that e is an internal edge with endpoints u and v, where u is white. Let e' and e'' be the other edges containing u and let A_e, $A_{e'}$ and $A_{e''}$ each be as above with A_e on the u side of e and $A_{e'}$, $A_{e''}$ on the non-u side of e' and e''. Then

$$\sum_{i\in A_e} x_e = \sum_{i\in A_{e'}} x_i + \sum_{A_{e''}} x_i \geq a_{e'} + a_{e''} = a_e.$$

A similar argument show that $\sum_{i\in A_e} x_i \geq a_e$ is redundant if v is black.

We now must show that each of these inequalities does define a facet. Our proof is by induction on n. First, consider the inequality $\sum_{i\in A_e} x_e \geq a_e$ where e connects a black vertex x to a white vertex y and A_e is on the x side of e. Let i and j be two leaves of T that border a common vertex v of T. Define T' and ϕ as in the proof of the previous proposition and let x' be a point of T' where all inequalities except for the one arising from the edge e in T are strict. Then any point x obeying $\phi(x) = x'$ and x_i, $x_j \in (0,1)$ will have all inequalities but the required one strict. Similar arguments apply to the other inequalities. $\square$

We thus see that the facets of $\Pi(T)$ which correspond to loop and co-loop free matroids are in bijection with the edges joining vertices of different colors, the facets corresponding to matroids with a loop correspond to edges joining leaves to black vertices and the facets corresponding to matroids with a co-loop correspond to edges from leaves to white vertices.

Let T be a trivalent tree. Suppose that e is an internal edge. Define $T\backslash\backslash e$ to be the pair of trivalent trees formed as follows: Let u and v be the endpoints of e and let $\{a, b, v\}$ and $\{c, d, u\}$ be the neighbors of u and v respectively. Delete the vertices u and v and all edges ending at them. Then draw new edges joining a to b and c to d. If e is an edge joining a leaf i to a vertex v, we define $T\backslash\backslash e$ similarly except that we simply delete the vertex i. We will use the same notation $F\backslash\backslash e$ for the analogous construction when F is a forest of trivalent trees and e an internal edge of a tree of F. If T is colored, we let $T\backslash\backslash e$ inherit the coloring of T in the obvious way.

Proposition 4.6.3. *Let e be an internal edge of T joining two vertices of opposite colors. Then the corresponding facet of $\Pi(T)$ is $\Pi(T_1) \times \Pi(T_2)$ where $T\backslash\backslash e = T_1 \sqcup T_2$. If e joins a leaf i to a vertex v then the corresponding facet of $\Pi(T)$ is $\Pi(T\backslash\backslash e) \times P_{\{i\}}$ where $\{i\}$ is given the structure of a rank 1 matroid if v is black and a rank 0 matroid if v is white.*

The faces of $\Pi(T)$ *are in bijection with the possible colored forests that can result from repeated splittings of* T *along edges connecting vertices of opposite color or connecting leaves to the rest of* T. *The faces of* $\Pi(T)$ *that correspond to loop and co-loop-free matroids are in bijection with the possible colored forests that can arise by repeated splittings along internal edges between vertices of opposite color.*

Proof. The first paragraph follows because it is easy to check that both polytopes are defined by the same inequalities. The second paragraph follows by repeatedly using the first. □

4.7 Special Cases of the f-Vector Conjecture

In this section we will prove the following results:

Theorem 4.7.1. *Let* L *be a* d*-dimensional tropical linear space in* n *space. Then* L *has at most* $\binom{n-2}{d-1}$ *vertices, with equality if and only if* L *is series-parallel.*

Theorem 4.7.2. *Let* L *be a* d*-dimensional tropical linear space in* n *space with* $n = 2d$ *or* $n = 2d+1$. *Then* L *has at most* 1 *bounded facet if* $n = 2d$ *and at most* d *if* $n = 2d + 1$.

We first recall the definition of the Tutte polynomial of a matroid. If M is a rank d matroid on the ground set $[n]$ and $Y \subseteq [n]$, let $\rho_M(Y)$ denote the rank of Y. The polynomial

$$r_M(x,y) = \sum_{Y \subseteq [n]} x^{|Y|-\rho_M(Y)} y^{d-\rho_M(Y)}$$

is known as the rank generating function of M. The polynomial

$$t_M(z,w) = r_M(z-1, w-1)$$

is known as the Tutte polynomial of M. Almost all matroid invariants can be computed in terms of the Tutte polynomial; see Chapter 6 of [50] for a survey of its importance. Write $t_M(z,w) = \sum t_{ij} z^i w^j$. Although not obvious from this definition, all of the t_{ij} are nonnegative. For $n \geq 2$ we have $t_{10} = t_{01}$, this number is known as the beta invariant of M and denoted $\beta(M)$.

We will need the following result:

Proposition 4.7.3. *Let M be a matroid on at least 2 elements. Then $\beta(M) = 0$ if and only if M is disconnected and $\beta(M) = 1$ if and only if M is series-parallel.*

Proof. See [10], theorem II, for the first statement and [8], theorem 7.6, for the second. □

The key to proving Theorem 4.7.1 will be proving the following formula:

Lemma 4.7.4. *Let M be a matroid and let $\mathcal{D}$ be a matroidal subdivision of P_M. Let $\mathring{\mathcal{D}}$ denote the set of internal faces of $\mathcal{D}$. Then*

$$t_M(z,w) = \sum_{P_\gamma \in \mathring{\mathcal{D}}} (-1)^{\dim(P_M)-\dim(P_\gamma)} t_\gamma(z,w).$$

Before proving Lemma 4.7.4, let us see why it implies Theorem 4.7.1. Considering the case where $M = \Delta(d, n)$ and comparing the coefficients of z on each side, we see that

$$\beta(\Delta(d,n)) = \binom{n-2}{d-1} = \sum_{\substack{\gamma \in \mathcal{D} \\ \gamma \text{ a facet}}} \beta(\gamma).$$

(Note that all nonfacets of $\mathring{\mathcal{D}}$ correspond to disconnected matroids and hence have beta invariant 0. and that every facet of $\mathcal{D}$ is in $\mathring{\mathcal{D}}$.)

Every term in the right hand sum is a positive integer. Thus, $\mathcal{D}$ has at most $\binom{n-2}{d-1}$ facets, with equality if and only if all of the facets have beta invariant 1, *i.e.*, if and only if $\mathcal{D}$ is a decomposition into series-parallel matroids. □

We now prove Lemma 4.7.4.

Proof. Since $t_M(z,w) = r_M(z-1, w-1)$, it is enough to prove

$$r_M(x,y) = \sum_{P_\gamma \in \mathring{\mathcal{D}}} (-1)^{\dim(P_M) - \dim(P_\gamma)} r_\gamma(x,y).$$

Plugging in the definition of r_M and interchanging summation, it is enough to show that for every $Y \subseteq [n]$,

$$x^{|Y| - \rho_M(Y)} y^{d - \rho_M(Y)} = \sum_{\gamma \in \mathring{\mathcal{D}}} (-1)^{\dim(P_M) - \dim(P_\gamma)} x^{|Y| - \rho_\gamma(Y)} y^{d - \rho_\gamma(Y)}.$$

Comparing coefficients of $x^{|Y|-r} y^{d-r}$, we are thus being asked to show that

$$\sum_{\substack{P_\gamma \in \mathring{\mathcal{D}} \\ \rho_\gamma(Y) = r}} (-1)^{\dim(P_M) - \dim(P_\gamma)} = \left\{ \begin{matrix} 1 \text{ if } r = \rho_M(Y) \\ \\ 0 \text{ if } r < \rho_M(Y) \end{matrix} \right\}.$$

The sum is empty if $r > \rho_M(Y)$. Equivalently, we will show

$$\sum_{\substack{P_\gamma \in \mathring{\mathcal{D}} \\ \rho_\gamma(Y) \geq r}} (-1)^{\dim(P_M) - \dim(P_\gamma)} = 1$$

for all $r \leq \rho_M(Y)$.

Let ℓ_Y be the linear function $\Delta(d,n) \to \mathbb{R}$ mapping $(x_i) \mapsto \sum_{i \in Y} x_i$. Then $\rho_\gamma(Y) = \max_{x \in P_\gamma} \ell_Y(x)$. Thus, we see that $\rho_\gamma(Y) \geq r$ if and only if γ has a nonempty intersection with the half space $\ell_Y > r - 1/2$. The promised equality now follows by the following lemma applied to the polytope $P_M \cap \{x : \ell_Y(x) > r - 1/2\}$.

Lemma 4.7.5. *Let P be any bounded polytope and Γ the internal faces of a decomposition of P. Then $\sum_{\gamma \in \Gamma} (-1)^{\dim(P) - \dim(\gamma)} = 1$.*

Proof. This sum is $(-1)^{\dim P}(\chi(P) - \chi(\partial P))$ where χ is the Euler characteristic. As P is contractible and ∂P is a sphere of dimension $\dim(P) - 1$, the result follows. $\square$

$\square$

One might hope to use the higher degree terms of Lemma 4.7.4 to produce additional bounds on the f-vector of $\mathcal{D}$. Unfortunately, Lemma 4.7.4 is incapable of producing a complete set of restrictions. For example, consider the matroids M_1, M_2 and M_3 corresponding to the graphs in Figure 4.13.

The Tutte polynomials of these matroids are

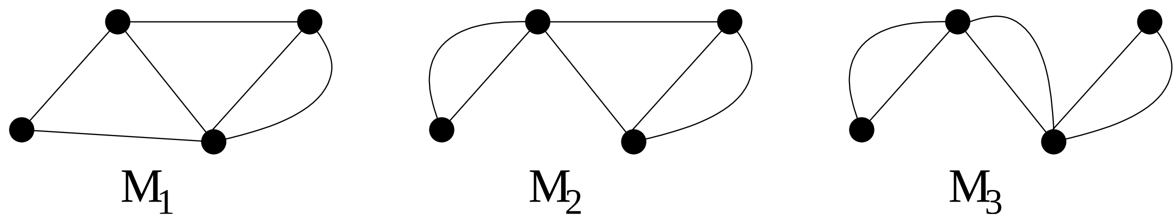

Figure 4.13: The Graphical Matroids M_1, M_2 and M_3

$$\begin{aligned} t_{M_1} &= z + w + 2z^2 + 3zw + 2w^2 + z^3 + z^2w + wz^2 + w^3 \\ t_{M_2} &= z^2 + 2zw + w^2 + z^3 + 2z^2w + 2zw^2 + w^3 \\ t_{M_3} &= z^3 + 3z^2w + 3zw^2 + w^3. \end{aligned}$$

We have

$$6t_{M_1} - 9t_{M_2} + 4t_{M_3} = t_{\Delta(3,6)} = 6z + 6w + 3z^2 + 3w^2 + z^3 + w^3.$$

Nonetheless, it follows from Theorem 4.7.2 (proved below) that there is no matroidal decomposition of $\Delta(3,6)$ which has 4 internal faces of codimension 2.

Proof of Theorem 4.7.2. Suppose that L is a d-dimensional tropical linear space in n-space. Let $M^\vee$ be a bounded facet of L, dual to a polytope P_M. Then the matroid M has rank d and has d connected components. So $M = \bigoplus_{i=1}^d M_i$ where each M_i has rank 1 and, because $M^\vee$ is bounded, $|M_i| \geq 2$. Let S_i be the ground set of M_i. The vertices of P_M are those subsets of $[n]$ that meet each S_i precisely once.

Now, consider two such facets $M^\vee$ and $(M')^\vee$ occurring in the same linear space and define S_i and S_i' as above. The polytope $P_M \cap P_{M'}$ must be either matroidal or empty.

Lemma 4.7.6. *Let* $\Gamma = \Gamma(M, M')$. Γ *may have multiple edges, let* $\overline{\Gamma}$ *be the simple graph obtained by replacing each multiple edge of* Γ *by a single edge. Then* $P_M \cap P_{M'}$ *is empty if and only if* $\overline{\Gamma}$ *has no perfect matchings.* $P_M \cap P_{M'}$ *is matroidal if and only if* $\overline{\Gamma}$ *has precisely one perfect matching.*

A *perfect matching* of a graph Γ is a collection of edges such that, for every vertex v of Γ, exactly one edge in the collection contains v.

Proof. The vertices of $P_M \cap P_{M'}$ are in bijection with the perfect matchings of Γ. There is an obvious surjection from perfect matchings of Γ to matchings of $\overline{\Gamma}$. This proves the first claim.

Suppose that $P_M \cap P_{M'}$ is nonempty and matroidal. Let (B_1, B_2) be an edge of $P_M \cap P_{M'}$, so B_1 and B_2 correspond to perfect matchings of Γ that differ only by a single edge say $B_1 \setminus \{e_1\} = B_2 \setminus \{e_2\}$. Then e_1 and e_2 have the same image in $\overline{\Gamma}$ in order for both B_1 and B_2 to be perfect matchings. Since every pair of vertices of $P_M \cap P_{M'}$ is connected by a chain of edges, we see that every perfect matching of Γ gives rise to the same perfect matching of $\overline{\Gamma}$ or, in other words, $\overline{\Gamma}$ has exactly one perfect matching. Conversely, if $\overline{\Gamma}$ has exactly one perfect matching, it is clear that $P_M \cap P_{M'}$ is a product of simplices and hence matroidal. $\square$

We now turn to the cases we are interested in, when $n = 2d$ or $2d + 1$. We maintain the notation of the lemma. First, suppose that $n = 2d$. Then every S_i must have order 2 so, for any S_i and S_i', the graph Γ must consist of disjoint cycles. The only way that $\overline{\Gamma}$ can be a forest is if all of those cycles have length 2. But then $\{S_i\} = \{T_i\}$. So we see that a linear space of dimension d in $2n$ space can have at most one bounded facet.

Now suppose that $n = 2d + 1$. Then all of the S_i have order 2 except for one which has order 3. We order the S_i so that $|S_1| = 3$. Γ consists of several disjoint cycles and one component C. C consists of a pair of vertices v and w of degree 3 with opposite colors and three disjoint paths γ_1, γ_2, γ_3. Either these paths all run from v to w or there is one path each running from v to v, from v to w and from w to w.

In any of these cases, Γ has at least one perfect matching. One can check that the only case where $\overline{\Gamma}$ has only one perfect matching is when all of the cycles of $\Gamma \setminus C$ have length 2 and C consists of a path from v to w and two cycles of length 2, one containing v and the other containing w. The second condition is equivalent to requiring that $S'_{i'} \subset S_1$ and $S_i \subset S'_1$ for some i and i'.

Consider all of the three element subsets of $[n]$ that occur as an S_1 for some bounded facet of L. We first note that, if $S_1 = S'_1$ then $S_i = S'_i$ for every i (after reordering) as otherwise Γ would have a cycle of length more than 2. So all of the S_1^r subsets are different.

Let f be the number of bounded facets of L with $(S_1^r, S_2^r, \dots, S_d^r)$ the partition of $[n]$ corresponding to the r^{th} facet of L for $1 \le r \le f$. We now build a graph whose vertices are the various S_1^r and the labels $\{1, \dots, n\}$. Let there be an edge between S_1^r and k if $k \in S_1^r$. This graph has $f+n = f+2d+1$ vertices and $3f$ edges. Thus, if $f > d$, this graph has as many edges as vertices and must contain a cycle, we will use this cycle to obtain a contradiction. Saying this graph has a cycle means that there exist $i_1, \dots, i_g, j_1, \dots, j_g \in [n]$ and facets corresponding to collections (S_s^r), $1 \le s \le d$, $1 \le r \le g$ such that $S_1^r = \{i_r, i_{r+1}, j_r\}$, where the index r is cyclic modulo g.

Now, for any r and r', there must be some $S_s^r \subset S_1^{r'}$. Consider first the case where $r' = r+1$. Then $\{j_{r+1}, i_{r+2}\}$ must be one of the S_s^r, as we already know $i_{r+1} \in S_1^r$. Now consider when $r' = r+2$. We get that $\{j_{r+2}, i_{r+3}\}$ must be one of the S_s^r, as i_{r+2}, the other member of $S_1^{r'}$, is already contained in $\{j_{r+1}, i_{r+2}\}$. Continuing in this manner, we get that $\{j_{r+k}, i_{r+k+1}\}$ is an S_s^r for every k. But, when $k = g-1$, this intersects S_1^r, contradicting that the S_i^r are distinct.

$\square$

The f-vector conjecture predicts that, if $n = 2d+e$, L should have at most $\binom{d+e-1}{e}$ bounded facets. In the cases $e = 0$ and $e = 1$, the graph Γ always has at least one matching and the requirement that $\overline{\Gamma}$ have exactly one is very restrictive. Once $e = 2$, it is possible for $\overline{\Gamma}$ to have no perfect matchings and the problem grows much harder.

Nonetheless, I suspect that this kind of reasoning, looking only at the facets of L and not at the smaller spaces that contain them, is adequate to resolve the f-vector conjecture for facets once $e \geq 2$. More specifically, I conjecture:

Conjecture 4.7.7. *Let $n = 2d + e$. It is not possible to find more than $\binom{d+e-1}{e}$ distinct partitions of $[n]$ into d disjoint subsets, $[n] = S_1^r \sqcup \cdots \sqcup S_d^r$ where each S_s^r has order at least 2 and, for every r and r', the graph $\overline{\Gamma}$ formed from $\{S_1^r, \ldots, S_d^r\}$ and $\{S_1^{r'}, \cdots, S_d^{r'}\}$ has at most one perfect matching.*

Remark: Consider the case where $d = 3$. In this case, there is an elegant way to describe the compatability condition between (S_1, S_2, S_3) and (S_1', S_2', S_3'). Two partitions of $[n]$ into triples (S_1, S_2, S_3) and (S_1', S_2', S_3') are compatible if and only if, after reordering the S_i and S_i', we have $S_1 \subseteq S_1'$ and $S_3 \supseteq S_3'$.

4.8 Results on Constructible Spaces

In this section, we will prove several of the previously stated results on constructible spaces. Our first aim is to prove Theorem 4.1.3: that every constructible space is series-parallel. Clearly, $L^{\perp}$ is series-parallel if L is, since the dual of a series-parallel matroid is series-parallel. So it is enough to show that, if L and L' are two series-parallel tropical linear spaces that meet transversely then $L \cap L'$ is series-parallel as well.

Let L and L' be two tropical linear spaces of dimensions d and d' in n space which meet transversely at a point w. Let $\mathcal{D}$ and $\mathcal{D}'$ be the appropriate subdivisions of $\Delta(d,n)$ and $\Delta(d',n)$. Let $M^\vee$ and $(M')^\vee$ be the faces of $\mathcal{D}^\vee$ and $(\mathcal{D}')^\vee$ containing w. Let $M = \bigoplus_{i \in I} M_i$ and $M' = \bigoplus_{i \in I'} M'_i$ be the connected components of M and M'. We claim that all of the M_i and M'_i are series-parallel. P_M will be a face of some facet $P_{\tilde{M}}$ of $\mathcal{D}$, so this follows from

Lemma 4.8.1. *Let $\tilde{M}$ be a series-parallel matroid and let P_M be a face of $P_{\tilde{M}}$, assume that M is loop-free. Let $M = \bigoplus M_i$ be the decomposition of M into connected components. Then each of the M_i are series-parallel.*

Proof. This follows from the description of the faces of $\tilde{M}$ in Lemma 4.6.3. (Note that a rank one matroid on a single element is series-parallel.) □

The result now follows from

Proposition 4.8.2. *Let $M = \bigoplus M_i$ and $M' = \bigoplus M'_i$ be transverse matroids on $[n]$ with all of the M_i and M'_i series-parallel. Then $M \wedge M'$ is a direct sum of series-parallel matroids, and is series-parallel if it is connected.*

Proof. $M \wedge M'$ is formed by a sequence of parallel connections. The parallel connection of two series-parallel matroids is series-parallel. □

Our next goal is to prove Theorem 4.1.2 – every constructible space achieves the f-vector of the f-vector conjecture. Our strategy is as follows: A constructible space is made by succesive dualization and transverse intersection.

It is easy to show that dualization preserves the f-vector, so our main goal is to understand the effect of transverse intersection.

We will first prove (Theorem 4.8.4) that, if L and L' are series-parallel tropical linear spaces and $v \in \mathbb{R}^n$, then the f-vector of $L \cap (L'+v)$ is independent of v as long as we impose that L and $L'+v$ meet transversely. The most technical part of this proof is the analysis of perturbing L' slightly when L and L' are almost transverse, this is carried out in Lemma 4.8.5. Our result ultimately relies on a lemma (Lemma 4.8.3) about chains of flats in series-parallel matroids.

Once we have proven that the f-vector of $L \cap (L' + v)$ is independent of v, we find a particular value of v for which this f vector is particularly easy to compute. The computation involves several binomial coefficient identities.

Lemma 4.8.3. *Let M be a series-parallel matroid. Call a flat Q of M a* flacet *if $M|_Q$ and M/Q are connected. Let f and $g \in M$. Then the collection of all flacets Q such that $f \in Q$ and $g \notin Q$ forms a chain $Q_1 \subset \cdots \subset Q_d$. Moreover, the length d of this chain is preserved when the roles of f and g are switched.*

We define $d(f, g)$ to be d.

Remark: The terminology "flacet" was suggested by Bernd Sturmfels. The motivation for this terminology is that, for M any connected matroid, the flacets of M are in bijection with the facets of P_M.

Proof. As remarked above, if Q is a flacet then $M|_Q \oplus M/Q$ corresponds to a facet of M. Let $M = \mu(T)$. By Lemma 4.6.3, all of the facets of P_M correspond to either edges of T which contain a leaf or edges of T that connect to vertices of opposite colors. Under that correspondence, an edge e corresponds to a flacet that contains f and not g if and only if e separates f and g with the end closer to g colored black and the end closer to f either colored white or equal to e. Consider the path γ through T connecting f and g, divide γ into alternating blocks of white and black vertices. Then we see that the number of flacets containing f and not g is equal to the number of black blocks. This number is the same when f and g are switched. The flat corresponding to e consists of the leaves of T that are on the non-black end of e. Thus, all of the flats in question are nested. □

Our first major goal on the way to proving Theorem 4.1.2, nontrivial in its own right, is to prove:

Theorem 4.8.4. *Let L and L' be two series-parallel tropical linear spaces in n-space, of dimensions d and d'. Let a and $b \in \mathbb{R}^n$ be such that L meets $L' + a$ and $L' + b$ transversely. Then $L \cap (L' + a)$ and $L \cap (L' + b)$ have the same bounded f-vector.*

We first prove a lemma describing what happens when we perturb a non-transverse intersection of two tropical linear spaces.

Let L and L' be tropical linear spaces of dimensions d and d'. Let 0 be a vertex of $L \underset{\text{stable}}{\cap} L'$ and let $M = \bigoplus_{i \in I} M_i$ and $M' = \bigoplus_{i' \in I'} M'_{i'}$ be the matroids such that 0 lies in the relative interiors of $M^\vee$. Let $\mathcal{E}$ be the subdivision of $\Delta(d + d', n)$ corresponding to $L \underset{\text{stable}}{\cap} L'$, so $P_{M \wedge M'}$ is a face of $\mathcal{E}$ with 0 in the relative interior of $(M \wedge M')^\vee$.

Let v be a generic vector of $\mathbb{R}^n$ and consider the tropical linear space $L \underset{\text{stable}}{\cap} (L'+v) = L \cap (L'+v)$. The tropical linear space $L \underset{\text{stable}}{\cap} (L'+v)$ corresponds to a matroidal subdivision $\mathcal{E}_v$ of $\Delta(d+d', n)$. We will describe the facets into which the relative interior of $P_{M \wedge M'}$ is divided in $\mathcal{E}_v$.

Let T be a tree and $\pi : T \to \Gamma(M, M')$ a map of graphs which is bijective on edges. (Note: since we assumed that $(M \wedge M')^\vee$ was a vertex of $L \underset{\text{stable}}{\cap} L'$, we know that $\Gamma(M, M')$ is connected.) Let u be a real valued function on the vertices of T all of whose values are distinct. Let i be a vertex of $\Gamma(M, M')$. Let $\{j_1, l \ldots, j_r\} = \pi^{-1}(i)$ ordered so that $u(j_1) < u(j_2) < \cdots < u(j_r)$.

Recall that $S_i^{(')}$ is the set of edges ending at vertex i and has the matroid structure $M_i^{(')}$ placed on it. Let Q_k be the subset of $S_i^{(')}$ consisting of those edges whose preimage in F has endpoint $j_{k'}$ for some $k' \leq k$. Thus $Q_1 \subset Q_2 \subset \cdots \subset Q_k$. Place the matroid structure $M_i^{(')}|_{Q_k}/Q_{k-1}$ on the edges of T coming into vertex j_k . (We adopt the convention that Q_0 is the empty set.)

Now T is a tree where, at every vertex, the structure of a matroid has been assigned to the edges ending at that vertex. We denote by $\mu(M, M', T, \pi, u)$ the matroid formed by taking parallel connections along the edges of T as in Section 4.4.

Lemma 4.8.5. *With the above notation, and for v sufficiently small, every facet of $\mathcal{E}_v$ contained in the relative interior of $P_{M \wedge M'}$ is of the form $\mu(M, M', T, \pi, u)$*

for some T, π and u. More specifically, $\mu(M, M', T, \pi, u)$ occurs in $\mathcal{E}$ if and only if

1. *For every vertex i of $\Gamma(M, M')$ and for every $j_k \in \pi^{-1}(i)$, the matroid $M_i^{(\prime)}|_{Q_k}/Q_{k-1}$ is connected.*

2. *For every edge $e = (i, i')$ of T with $\pi(i) \in I$ and $\pi(i') \in I'$, we have $v(\pi(e)) = u(i) - u(i')$. Here $v \in \mathbb{R}^n$ is thought of as a function on the edges of $\Gamma(M, M')$.*

More generally, suppose that $(M \wedge M')^\vee$ is an i-dimensional face of $L \underset{stable}{\cap} L'$ (equivalently, if $\Gamma(M, M')$ has i connected components). Then $P_{M \wedge M'}$ is subdivided in $\mathcal{E}_v$ into facets. These facets are described by giving a forest F, a map $\pi : F \to \Gamma(M, M')$ in which the preimage of each component of $\Gamma(M, M')$ is a tree and a function u on the vertices of F which obeys the hypotheses above. We use the notation $\mu(M, M', F, \pi, u)$ in this situation.

Remark: The assumption that $M_i^{(\prime)}|_{Q_k}/Q_{k-1}$ is loop-free is equivalent to assuming that Q_{k-1} is a flat of $M_i^{(\prime)}|_{Q_k}$. Assuming that $M_i^{(\prime)}|_{Q_k}/Q_{k-1}$ is loop-free for all k is thus equivalent to assuming that $Q_1 \subset Q_2 \subset \cdots \subset Q_k$ is a chain of flats of $M_i^{(\prime)}$.

Proof. Every face of $L \cap (L' + v)$ is the intersection of a face $R^\vee$ of L and a face $(R')^\vee + v$ of $L' + v$. This matroid corresponding to said face of $L \cap (L' + v)$ is $R \wedge R'$. If this face is dual to one of the face appearing in the subdivision of the relative interior of $P_{M \wedge M'}$ then $M^\vee \subset R^\vee$ and $(M')^\vee \subset (R')^\vee$. That $M^\vee \subseteq R^\vee$ means that R is of the form $\bigoplus Q_k/Q_{k-1}$ for $Q_\cdot$ a chain of flats in M. Similarly, $R' = \bigoplus Q'_k/Q'_{k-1}$ for $Q'_\cdot$ a chain of flats in M'. By refining the chains $Q_\cdot^{(')}$, we may assume that each succesive quotient $Q_k^{(')}/Q_{k-1}^{(')}$ is connected.

Let T be the graph $\Gamma(R, R')$. Since we are looking for facets of $\mathcal{E}_v$, we have that $R \wedge R'$ is connected so T is connected. The fact that the faces $R^\vee$ and $(R')^\vee + v$ meet for generic v means that T is acyclic. Thus, T is a tree. We have a map $\pi : \Gamma(R, R') \to \Gamma(M, M')$ which sends $Q_k^{(')}/Q_{k-1}^{(')}$ to the unique connected component of $M^{(')}$ which contains all the members of the set $Q_k^{(')} \setminus Q_{k-1}^{(')}$. This map is bijective on edges. (Both edge sets are naturally labelled with $[n]$ and this map preserves the labelling.) Find a real valued function u on the vertices of T such that u orders the vertices above each $i \in I \cup I'$ correctly and we will have $\mu(M, M', T, \pi, u) = R \wedge R'$.

Now, not every $\mu(M, M', T, \pi, u)$ occurs because sometimes the relative interiors of $R^\vee$ and $(R')^\vee + v$ are disjoint. We must determine when the relative interiors of $R^\vee$ and $(R')^\vee + v$ meet. For v small enough, this is equivalent to determining when v is in the relative interior of the sum of the local cones at 0 of $R^\vee$ and $-(R')^\vee$. The local cone of R is spanned by all vectors $x \in \mathbb{R}^n$ with x constant on the set $Q_k \setminus Q_{k-1}$ and $x(Q_k \setminus Q_{k-1}) < x(Q_{k+1} \setminus Q_k)$. (See Proposition 4.2.3.) Such a vector x gives rise to a function u on $\pi^{-1}(I)$, where $u(j) = x(e)$ for any edge e ending at j. Similarly, we can use a point of the local cone of $-(R')^\vee$

to define a function u on $\pi^{-1}(I')$. We thus see that the condition that v lie in the relative interior of the sum of the local cones of $R^\vee$ and $-(R')^\vee$ implies that the function u obeys condition (2) above. We have thus shown that every facet into which the relative interior of $P_{M\wedge M'}$ is divided arises in the above manner.

It is easy to reverse this argument and show that if T, π and u meet the conditions above then the faces $R^\vee$ and $R^\vee + v$ do meet. The face given by the intersection is dual to the matroid $R \wedge R'$, which is formed by contracting along the edges of T (see Lemma 4.4.3). As v approaches 0, the vertex $R^\vee \cap (R^\vee + v)$ approaches $(M \wedge M^\vee)$. This tells us that $P_{R\wedge R'}$ appears in the relative interior of $P_{M\wedge M'}$.

There is no difficulty except for additional bookkeeping in the case where $M \wedge M'$ has multiple components. $\square$

Remark: I haven't found a comparably simple dsescription of the lower dimensional faces into which the relative interior $P_{M\wedge M'}$ is dubdivided. The most obvious guess is to replace the tree T with a forest. This winds up describing the faecs of $\mathcal{E}_v$ contained in $P_{M\wedge M'}$, but it describes both the faces in the relative interior of $P_{M\wedge M'}$ ans the boundary. The trouble is that, as v approaches 0, $R^\vee \cap (R^\vee + v)$ may shrink down to the vertex $(M \wedge M')^\vee$ or it may remain higher dimensional and approach a face which contains $(M \wedge M')^\vee$ in its boundary. In the setting of the lemma, $R^\vee \cap (R^\vee + v)$ is already a vertex, so this issue doesn't arise.

We now prove Theorem 4.8.4.

Proof of Theorem 4.8.4. Let $U_i \in \mathbb{R}^n$ be the set of v such that, for any two faces F and F' of L and L', either the relative interiors of F and $F' + v$ are disjoint or they span an affine linear space of dimension at least $n - i$. Thus, L and $L' + v$ meet transversely if and only if $v \in U_0$. $\mathbb{R}^n \setminus U_i$ is a polyhedral complex of dimension at most $n - i - 1$. In particular, we see that U_1 is path connected. We may therefore join a and b by a path through U_1. It is thus enough to show that, if $v \in U_1 \setminus U_0$ and we perturb v into U_0 then the f-vector we get is independent of the choice of perturbation. Without loss of generality, we may assume that $v = 0$.

Let $N^\vee$ be a face of $L \underset{\text{stable}}{\cap} L^\vee$ at which two faces $M^\vee$ and $(M')^\vee$ fail to meet transversely. Let the connected components of M and M' be $\bigoplus_{i \in I} M_i$ and $\bigoplus_{i \in I'} M'_i$.

In our setting, since $M^\vee$ and $(M')^\vee$ span an affine hyperplane, $\Gamma(M, M')$ will have first Betti number 1 and will thus contain a unique cycle. Let $e_1, \dots, e_{2r}$ be the edges of that cycle in cyclic order. Write i_j and i_{j+1} for the ends of e_j with $i_j \in I$ for j odd and $i_j \in I'$ for j even. Let G be the component of $\Gamma(M, M')$ containing this cycle.

Now, let $v \in \mathbb{R}^n$ be generic and sufficiently small and let $\mathcal{E}_v$ be the decomposition of $\Delta(d+d'-n, n)$ associated to $L \underset{\text{stable}}{\cap} (L'+v)$. $\mathcal{E}_{\epsilon v}$ is a subdivision of $\mathcal{E}_0$. We will show that the number of faces of codimension c into which the relative interior of P_N is subdivided is independent of v. We will first compute this number on the assumption that $\sum_{s=1}^{2r}(-1)^s v_{e_s} > 0$. We will then recompute it on the assumption that $\sum_{s=1}^{2r}(-1)^s v_{e_s} < 0$ and see that we get the same answer.

We first consider the facets into which P_N is subdivided. In the notation of Lemma 4.8.5, every face of $\mathcal{E}_v$ is of the form $\mu(M, M', F, \pi, u)$. The map $\pi : F \to \Gamma(M, M')$ must be an isomorphism on every component of $\Gamma(M, M')$ except G and must be bijective on G except for a single vertex i_s somewhere in the loop of G whose preimage is two vertices. Let j_1 and j_2 be the two preimages of i_s, with the preimage of edge e_{s-1} ending at j_1 and the preimage of e_s ending at j_2.

Suppose that s is odd, so $i_s \in I$. Then condition (2) of Lemma 4.8.5 gives us

$$\begin{aligned} u(j_1) - u(j_2) &= \big(u(j_1) - u(\pi^{-1}(i_{s-1}))\big) + \big(u(\pi^{-1}(i_{s-1})) - u(\pi^{-1}(i_{s-2}))\big) + \cdots \\ &\qquad + \big(u(\pi^{-1}(i_{s+1})) - u(j_2)\big) \\ &= v(e_{s-1}) - v(e_{s-2}) + \cdots - v(e_s) \\ &= \sum_{s=1}^{2r}(-1)^s v_{e_s} > 0. \end{aligned}$$

If s is even, so $i_s \in I'$, then we have

$$\begin{aligned} u(j_1) - u(j_2) &= \big(u(j_1) - u(\pi^{-1}(i_{s-1}))\big) + \big(u(\pi^{-1}(i_{s-1})) - u(\pi^{-1}(i_{s-2}))\big) + \cdots \\ &\qquad + \big(u(\pi^{-1}(i_{s+1})) - u(j_2)\big) \\ &= -v(e_{s-1}) + v(e_{s-2}) - \cdots + v(e_s) \\ &= \sum_{s=1}^{2r} (-1)^s v_{e_s} > 0. \end{aligned}$$

Either way, $u(j_1) > u(j_2)$. So, in order to be in accord with condition (1) of Lemma 4.8.5, we must have that the edges coming into j_1 must form a flat Q of $M_{i_s}^{(')}$ and $M_{i_s}^{(')}|_Q$ and $M_{i_s}^{(')}/Q$ must both be connected. In other words, Q must be a flacet of $M_{i_s}^{(')}/Q$. For each flacet of $M_{i_s}^{(')}$ which contains e_{s-1} and not e_s, we get a facet of $\mathcal{E}_v$. Thus, we see that the number of facets of $\mathcal{E}_v$ into which P_N is subdivided is $D := \sum_{s=1}^{2r} d(e_{s-1}, e_s)$. If we redo this computation on the assumption that $\sum_{s=1}^{2r}(-1)^s v_{e_s} < 0$, we see that the number of facets into which P_N is subdivided is $\sum_{s=1}^{2r} d(e_s, e_{s-1})$ which, by Lemma 4.8.3, is the same as $\sum_{s=1}^{2r} d(e_{s-1}, e_s)$.

We now consider instead the problem of determining the number of codimension c faces into which the relative interior of P_N is subdivided in $\mathcal{E}_v$. Each such face is an intersection of facets of $\mathcal{E}_v$ in the relative interior of P_N. Once again, we start our discussion on the assumption that $\sum_{s=1}^{2r}(-1)^s v_{e_s} > 0$. We claim that every $c+1$ element subset of the D facets of $\mathcal{E}_v$ found above has a distinct, nonempty intersection in the interior of P_N. Thus, there are $\binom{D}{c+1}$ codimension c faces and this number would be the same if $\sum_{s=1}^{2r}(-1)^s v_{e_s} < 0$.

To give a facet of $\mathcal{E}_v$ is to give a choice of an integer s between 1 and $2r$ and a flacet Q of $M_{i_s}^{(')}$ which contains e_{s-1} and not e_s. To give $c+1$ such facets we must give, for each s a collection $Q_s^1, \ldots, Q_{2r}^{c_{2r}}$ of flacets with each Q_s^t containing e_{s-1} and not e_s, where $\sum c_s = c+1$. By Lemma 4.8.3, $Q_s^1, \ldots, Q_{2r}^{c_{2r}}$ will form a nested chain and we may reorder them such that $Q_s^1 \subset \cdots \subset Q_{2r}^{c_{2r}}$. We can find a c-tree forest F with a surjection π onto $\Gamma(M)$ where i_s has c_s preimages and the edges coming into the t^{th} preimage of i_s are labelled by the elements of the set $Q_s^t \setminus Q_s^{t-1}$. Let S be the matroid formed by making parallel connections along the edges of F. All of the facets meet at P_S. One can check that P_S meets the relative interior of P_N. The same argument goes through with e_s and e_{s-1} reversed when $\sum_{s=1}^{2r}(-1)^s v_{e_s} < 0$. $\square$

Remark: As v goes to 0, $(L \cap (L' + v))/(1, \ldots, 1)$ has a $(D-1)$-dimensional simplex which shrinks down to a point at $v = 0$. After $\sum_{s=1}^{2r}(-1)^s v_{e_s}$ changes sign, a new $(D-1)$-dimensional simplex begins to grow.

We now are ready to begin our final assault on Theorem 4.1.2. We will first need to study how the various faces of $\Delta(d, n)$ are subdivided in a constructible decomposition. Let S and T be disjoint subsets of $[n]$. Write $\Delta(d, n) \setminus S/T$ for the face of $\Delta(d, n)$ given by the equations $x_i = 0$ for $i \in S$ and $x_i = 1$ for $i \in T$. If p is a tropical Plücker vector then define $p \setminus S/T$ to be the function on the vertices of $\Delta(d - |T|, n - |S| - |T|)$ given by $(p \setminus S/T)_I = p_{I \cup T}$. It is clear that $p \setminus S/T$ obeys the tropical Plücker relations, as they are a subset of the tropical Plücker relations for p. So $p \setminus S/T$ corresponds to a tropical $d - |T|$ plane in $n - |S| - |T|$ space. If L is the tropical linear space corresponding to p then we write $L \setminus S/T$ for the tropical linear space corresponding to $p \setminus S/T$.

Lemma 4.8.6. *If L is a constructible d-plane in n-space and S and T are disjoint subsets of $[n]$ then $L \setminus S/T$ is also constructible.*

Proof. If L is a hyperplane then every $L \setminus S/T$ is either a hyperplane or the whole space in which it sits, and hence is constructible. Now suppose that the lemma holds for L. Then it holds for $L^{\perp}$ as $L^{\perp} \setminus S/T = (L \setminus T/S)^{\perp}$. Suppose that the lemma holds for L and L'. Then it holds for $L \underset{\text{stable}}{\cap} L'$ as $(L \underset{\text{stable}}{\cap} L') \setminus S/T = (L \setminus \emptyset/(S \cup T)) \underset{\text{stable}}{\cap} (L' \setminus S/T)$. □

Proof of Theorem 4.1.2. Our proof is by induction, first on n and then on the number of steps used to construct L. We set

$$f_{i,d,n} = \binom{n-2i}{d-i}\binom{n-i-1}{i-1} = \frac{(n-i-1)!}{(i-1)!(d-i)!(n-d-i)!}.$$

First, suppose that the theorem is true for L. $L^{\perp}$ and L have the same bounded f-vector, so the theorem for $L^{\perp}$ follows from the symmetry $f_{i,d,n} = f_{i,n-d,n}$. For the rest of the proof, suppose that L and L'' are a constructible d-plane and d' plane in n space which meet transversely and for which we know the theorem to hold. In addition, we know that the theorem holds for every $L \setminus S/T$ and $L'' \setminus S/T$ by our inductive hypothesis. Our aim will be to prove the theorem for $L \cap L''$.

By Theorem 4.8.4, it is enough to prove the result for $L \cap (L''+w)$ where w is chosen generically enough. We will take a w for which $w_1 \ll w_2 \ll \cdots \ll w_n$ and otherwise generic. Set $L' = L'' + w$.

Suppose that $M^\vee$ is a face of L dual to $P_M \subset \Delta(d,n)$. Then M is loop-free, let T be the set of co-loops of M so that M lies in the relative interior of $\Delta(d,n) \backslash \emptyset / T$. Then $M^\vee/(1,\ldots,1)$ is the product of $\mathbb{R}_+^T$ and a bounded polytope. Let $(M'')^\vee$ be a face of L'' and define T' similarly. We write $(M')^\vee = (M'')^\vee + w$ and $M' = M''$.

Define an ordered pair (T, T') of subsets of $[n]$ to be *nice* if, for all $1 \leq i < j \leq n$, either $j \in T$ or $i \in T'$.

Lemma 4.8.7. *For $w_1 \ll w_2 \ll \cdots \ll w_n$, we have $M^\vee \cap (M')^\vee \neq \emptyset$ if and only if (T,T') is nice. If (T,T') is nice then $M^\vee \cap (M')^\vee$ is bounded if and only if $T \cap T' = \emptyset$.*

.

Proof. First, suppose that (T,T') is not nice, so there exists some $i < j$ with $j \notin T$ and $i \notin T'$. On $M^\vee$, the function $x_j - x_i$ is bounded above (as $j \notin T$). On $(M'')^\vee$, $x_j - x_i$ is similarly bounded below. The value of $x_j - x_i$ on $(M'')^\vee$ is $w_j - w_i$ larger than on $(M'')^\vee$, so if $w_j - w_i$ is large enough then the values of this functional on $M^\vee$ and $(M'')^\vee + w$ will be distinct. Thus, $M^\vee$ and $(M'')^\vee + w$ will be disjoint.

Now suppose that (T,T') is nice. Let $(x_i) \in M^\vee$ and $(x_i'') \in (M'')^\vee$, set

$x'_i = x''_i + w_i$ and set $d_i = x'_i - x_i$. For w large enough, we have $d_1 < d_2 < \cdots < d_n$. The hypothesis that (T, T') is good implies that there is some $b \in [n]$ such that $i \in T'$ for every $i < b$ and $j \in T'$ for every $j > b$. After translating (x_i) by $(1, \ldots, 1)$, which will not change the assumption that $(x_i) \in M$, we may assume that $x_b = x'_b$. Consider the point

$$z := (x_1, \ldots, x_{b-1}, x_b, x_{b+1} + d_{b+1}, \ldots x_n + d_n)$$

$$(x'_1 - d_1, \ldots, x'_{b-1} - d_{b-1}, x'_b, x'_{b+1}, \ldots x'_n).$$

Clearly, $z \in (x_i) + \mathbb{R}^T_{\geq 0}$ and $z \in (x'_i) + \mathbb{R}^{T'}_{\geq 0}$, so $z \in M^\vee \cap (M')^\vee$. We have now proved the first claim.

Now assume that (T, T') is nice, so there is some point $z = (z_i) \in M^\vee \cap (M')^\vee$. Suppose that $i \in T \cap T'$. Then $z + ue_i \in M^\vee \cap (M')^\vee$ for every $u > 0$ and $M^\vee \cap (M')^\vee$ is unbounded.

Suppose, on the other hand, that $T \cap T' = \emptyset$. Consider any $i < j$. If neither i nor j is in T then $x_i - x_j$ is bounded on $M^\vee$; if neither i nor j is in T' then $x_i - x_j$ is bounded on $(M')^\vee$; if $j \in T$ and $i \in T'$ then $x_i - x_j$ is bounded above on $M^\vee$ and below on $(M')^\vee$. In any of these cases, we see that $x_i - x_j$ is bounded on $M^\vee \cap (M')^\vee$. So all of the $x_i - x_j$ are bounded on $M^\vee \cap (M')^\vee$ which implies that $(M^\vee \cap (M')^\vee)/(1, \ldots, 1)$ is bounded. $\square$

$M^\vee$ corresponds to a bounded face of $L \setminus \emptyset / T$. Writing $t = |T|$, if this bounded face is j-dimensional then $M^\vee$ is $(j + t)$-dimensional. Similarly writing $t' = |T'|$ and j' for the dimension of the face of $L' \setminus \emptyset / T'$ corresponding to $(M')^\vee$, the dimension of $(M')^\vee$ is $j' + t'$. The dimension of $M^\vee \cap (M')^\vee$ is then $j + j' + t + t' - n$. We thus see that the number of i-dimensional bounded faces of $L \cap L'$ is

$$\sum_{\substack{(T,T') \text{ nice} \\ T \cap T' = \emptyset}} \sum_{j+j'+t+t'-n=i} f_{j,d-t,n-t} f_{j',d'-t',n-t'}.$$

We have used our inductive assumption that all of the tropical linear spaces $L \setminus \emptyset / T$ and $L' \setminus \emptyset / T'$ have the correct f-vector.

The only nice pairs (T, T') with $T \cap T' = \emptyset$ are $(\{1, \ldots, b-1\}, \{b+1, \ldots, n\})$ and $(\{1, \ldots, b\}, \{b+1, \ldots, n\})$. So we can write this sum as

$$\sum_b \left(\sum_{j+j'=i+1} f_{j,d-b+1,n-b+1} f_{j',d'-n+b,b} + \sum_{j+j'=i} f_{j,d-b,n-b} f_{j',d'-n+b,b} \right).$$

We will now show that this sum is $f_{i,D,n}$, where D is $d + d' - n$.

We work with the first sum first. Plugging in the definition $f_{i,d,n} = \binom{n-d-1}{i-1}\binom{n-i-1}{n-d-1}$ and exchanging order of summation, we have

$$\sum_{j+j'=i+1} \binom{n-d-1}{j-1}\binom{n-d'-1}{j'-1} \sum_b \binom{n-b-j}{n-d-1}\binom{b-j'-1}{n-d'-1}.$$

Using the identity $\sum_{k+k'=p} \binom{k}{q}\binom{k'}{r} = \binom{p+1}{q+r+1}$, the inner sum evaluates to $\binom{n-i-1}{2n-d-d'-1} = \binom{n-i-1}{n-D-1}$. We are reduced to

$$\binom{n-i-1}{n-D-1} \sum_{j+j'=i+1} \binom{n-d-1}{j-1}\binom{n-d'-1}{j'-1}.$$

The remaining sum can be evaluated by the identity $\sum_{j+j'=r} \binom{p}{j}\binom{q}{j'} = \binom{p+q}{r}$. We get $\binom{n-i-1}{D-i}\binom{n-D-2}{i-1}$.

A similar argument shows that the second sum is $\binom{n-i-1}{D-i}\binom{n-D-2}{i-2}$. Adding these together, we get

$$\binom{n-i-1}{D-i}\left[\binom{n-D-2}{i-1}+\binom{n-D-2}{i-2}\right]=$$
$$\binom{n-i-1}{D-i}\binom{n-D-1}{i-1}=f_{i,D,n}.$$

□

4.9 Tree Tropical Linear Spaces

In this section, we will describe some tropical linear spaces which achieve the bounds of the f-vector conjecture and have very elegant combinatorics. Suppose that p_{ij} obeys the tropical Plücker relations so by the results in the preceding chapter p_{ij} corresponds to a tropical 2-plane L and to a tree T. We define a real valued function $\tau^d(p)$ on $\binom{[n]}{d}$ by $\tau^d(p)_I = \sum_{\substack{i,j\in I\\ i<j}} p_{ij}$.

Proposition 4.9.1. $\tau^d(p)$ *obeys the tropical Plücker relations.*

This can be checked in a routine manner, but we give a different proof that explains the motivation behind the formula. We will call the tropical linear space associated to $\tau^d(p)$ the d^{th} *tree space* of L and denote it $\tau^d(L)$.

Proof. We know that p_{ij} is realizable, meaning that we can find $x_1, \ldots x_n, y_1,$

$\ldots, y_n \in K$ such that $p_{ij} = v(x_i y_j - x_j y_i)$. Consider the linear space

$$\mathrm{RowSpan} \begin{pmatrix} x_1^{d-1} & x_2^{d-1} & x_3^{d-1} & \cdots & x_n^{d-1} \\ x_1^{d-2}y_1 & x_2^{d-2}y_2 & x_3^{d-2}y_3 & \cdots & x_n^{d-2}y_n \\ \vdots & \vdots & \vdots & \cdots & \vdots \\ y_1^{d-1} & y_2^{d-1} & y_3^{d-1} & \cdots & y_n^{d-1} \end{pmatrix}$$

The maximal minor $P_{i_1,\ldots,i_d}$ of this matrix is the Vandermonde determinant $\det(x_{i_r}^{d-s} y_{i_r}^{s-1}) = \prod_{i_r < i_s} (x_{i_r} y_{i_s} - x_{i_s} y_{i_r})$. So

$$v(P_I) = \sum_{\substack{i,j \in I \\ i<j}} v(x_{i_r} y_{i_s} - x_{i_s} y_{i_r}) = \sum_{\substack{i,j \in I \\ i<j}} p_{ij} = \tau^d(p)_I.$$

So the $\tau^d(p)_I$ come from an actual linear space and hence obey the tropical Plücker relations. $\square$

Remark: This construction is reminiscent of the construction of cyclic polytopes (see, for example, [51] example 0.6) and osculating flags (see [36], section 5), two other objects which realize maximal combinatorics.

In this section, we will give a complete description of the bounded part of $\tau^d(L)$ in terms of the combinatorics of T. We will make heavy use of the results of Section 4.6. From now on, we assume that T is trivalent.

Theorem 4.9.2. *The vertices of $\tau^d(L)$ are of the form $\mu(T, c)^\vee$ where c ranges over the $\binom{n-2}{d-1}$ ways to color the internal vertices of T black and white with $d-1$ colored black and $n-d-1$ colored white. The bounded i-dimensional faces of $\tau^d(T)$ are in bijection with the ordered pairs (F, c), where F is a forest with i*

trees that can be obtained by splitting T along internal edges and c is a black and white coloring of the internal vertices of F using $d-i$ black vertices and $n-d-i$ white vertices. (F,c) is contained in (F',c') if and only if (F',c') can be obtained from (F,c) by repeated splitting along edges connecting vertices of opposite colors.

Proof. It is enough to prove the first claim, as the rest then follows from the description of the faces of $\Pi(T,c)$ in Lemma 4.6.3. Let $\mathcal{D}$ be the subdivision of $\Delta(d,n)$ corresponding to $\tau^d(L)$, we must describe the facets of $\mathcal{D}$. If e is an edge of T, write $[n] = A_e \sqcup B_e$ for the partition of the leaves of T induced by splitting along e. If we have fixed a coloring of T, let a_e be the number of black vertices on the A_e side of e.

We have

$$p_I = \sum_{\substack{i,j\in I\\ i<j}} p_{ij} = (-1/2) \sum_{\substack{i,j\in I\\ i<j}} \sum_{e \text{ separates } i \text{ and } j} \ell(e) =$$

$$(-1/2) \sum_{e} \ell(e)|A_e \cap I| \left(d - |A_e \cap I|\right).$$

For any edge e of T, set $f_e(I) = -(1/2)\ell(e)|A_e \cap I| \left(d - |A_e \cap I|\right)$. For e containing a leaf, f_e is a linear functional. For e internal, f_e is convex (this uses $\ell(e) > 0$). So $\mathcal{D}$ is the common refinement of the subdivisions of $\Delta(d,n)$ induced by each of the convex functions f_e. The subdivision induced by f_e cuts $\Delta(d,n)$ into the pieces $k < \sum_{i\in A_e} x_i < k+1$ for $k \in \mathbb{Z}$.

So the facets of $\mathcal{D}$ are the nonempty sets of the form

$$\left\{x : k_e < \sum_{i \in A_e} x_i < k_e + 1 \ \forall_{e \in T}\right\}$$

for some integers k_e. When the k_e arise as the a_e for some coloring c, this set is the interior of $\Pi(T, c)$ and hence not empty. We now must show that all the facets of $\mathcal{D}$ arise in this manner. It is enough to show that a generic point of $\Delta(d, n)$ lies in $\Pi(T, c)$ for some coloring c of T.

Let x_i be a generic point of $\Delta(d, n)$ and let k_e and l_e be the integers such that $k_e < \sum_{i \in A_e} x_i < k_e + 1$ and $l_e < \sum_{i \in B_e} x_i < l_e + 1$, so $k_e + l_e = d - 1$. Let v be an internal vertex of T with edges e_1, e_2 and e_3 ending at v. Without loss of generality, suppose that A_{e_s} is the non-u side of e_s for $s = 1, 2, 3$. Then

$$k_{e_1} + k_{e_2} + k_{e_3} < \sum_{s=1}^{3} \sum_{i \in A_{e_s}} x_i = d < k_1 + k_2 + k_3 + 3$$

so $k_{e_1} + k_{e_2} + k_{e_3} = d - 1$ or $d - 2$. We color v white if this sum is $d - 1$ and black if it is $d - 2$.

We claim that precisely $d - 1$ vertices are colored black. Let B denote the set of vertices that are colored black and W those that are colored white. We have

$$\sum_{e \in T} (k_e + l_e) = (2n - 3)(d - 1) = |B|(d - 2) + |W|(d - 1) + n(d - 1)$$

where the second expression comes from grouping the sum on vertices. We have $|B| + |W| = n - 2$ so

$$|B|(d - 2) + |W|(d - 1) + n(d - 1) = (d - 1)(2n - 2) - |B|.$$

We deduce that $(2n-3)(d-1) = (2n-2)(d-1) - |B|$ and thus $|B| = d-1$. We now will show that (x_i) is contained in the polytope $\Pi(T, c)$.

Let e be an internal edge of T and let a_e be the number of black vertices on a chosen side of e. Then $k_e < \sum_{i \in A_e} x_i < k_e + 1$ and we want to show $a_e < \sum_{i \in A_e} x_i < a_e + 1$. Let S be the subtree of T lying on the A_e side of e and let S have s vertices, $s - |A_e|$ internal vertices and $s-1$ edges. We have

$$\sum_{e \subset S} (k_e + l_e) = (s-1)(d-1) = a_e(d-2) + (s - |A_e| - a_e)(d-1) + |A_e|(d-1) - l_e$$

by once again grouping the sum on edges and on vertices. Canceling, we get $a_e = d - 1 - l_e = k_e$.

So we have shown that a generic point of $\Delta(d, n)$ lies in $\Pi(T, c)$ for some coloring c of T and we are done. □

Corollary 4.9.3. *$\tau^d(L)$ has $f_{i,d,n} = \binom{n-2i}{d-i}\binom{n-i-1}{i-1}$ bounded faces of dimension i.*

Proof. We must count the i tree forests F which can be obtained by repeated splittings of T and then multiply this number by $\binom{n-2i}{d-i}$, the number of ways to choose which vertices to color black. Thus, it is enough to show that every trivalent tree T can be split into $\binom{n-i-1}{i-1}$ different i tree forests by splitting along internal edges. Let $F_{i,n}$ denote the number of ways to split a trivalent n leaf tree into an i tree forest. We will prove $F_{i,n} = \binom{n-i-1}{i-1}$ by induction on n, it is clearly true for $n = 3$.

Let a and b be two leaves of T that have a common neighbor v. Let F be an i tree forest obtained by splitting T along internal edges. Clearly, no sequence of splittings can ever separate a and b. There are then two cases: the tree of F that contains $\{a, b\}$ contains no other leaves, or it has some other leaf.

Case I: $\{a, b\}$ is a component of F. Let e be the edge joining v to the rest of T. Then, at some point in the splitting procedure, e is split and we may as well assume that it is the first step. Let T' be the component of $T\backslash\backslash e$ other than $\{a, b\}$. The number of F for which this case applies is then the number of $i-1$ tree forests obtainable by splitting T', or $F_{i-1,n-2}$.

Case II: The component of F containing $\{a, b\}$ has additional vertices. Let T'' be the tree obtained by shrinking a, b and v down to a single leaf. The number of F for which this case applies is the same as the number of i tree forests obtainable by splitting T'', or $F_{i,n-1}$.

So $F_{i,n} = F_{i-1,n-2} + F_{i,n-1} = \binom{n-i-2}{i-2} + \binom{n-i-2}{i-1} = \binom{n-i-1}{i-1}$ as desired. $\square$

Chapter 5

Tropical Curves

In this chapter we will attempt to describe which subsets of $\mathbb{R}^n$ can occur as $\operatorname{Trop} X$ for X a curve in $\mathbb{R}^n$. By Theorem 2.4.5, $\operatorname{Trop} X$ will be a 1-dimensional polyhedral complex, which is to say, a graph. This graph must obey the zero tension condition (Theorem 2.5.1).

If such a curve X exists then, as described in Section 2.4, the tropical degeneration of X will be a curve X_0 which can be broken up into components indexed by the vertices of $\operatorname{Trop} X$ whose intersections are indexed by the edges of $\operatorname{Trop} X$. There is an unfortunate issue that these components themselves may be disconnected but, as noted in Proposition 2.5.2, we can often rule that possibility out by the zero tension condition. In this case the first Betti number of $\operatorname{Trop} X$ is a lower bound for the genus of X. We will concentrate our attention on the case where the first Betti number of $\operatorname{Trop} X$ is equal to the genus of X. In that case, every component of X_0 must be rational and we will be able to apply ideas of Mumford to the problem.

The zero tension condition is not sufficient to guarantee the existence of X of appropriate genus and degree. Consider the following example of Mikhalkin: Let C be the zero tension curve in three space shown in Figure 5.1. (In this figure, everything to the left of the points P, Q and R is in a two-dimensional plane.) Note that we can preserves the abstract graph Γ and the slopes of the edges while varying C in a 13 dimensional family: we can translate C within $\mathbb{R}^3$ (3 dimensions), vary the side lengths of the hexagon (4 dimensions, since it must remain a closed loop) and vary the positions of points P, Q, R, S, T and U along lines (1 dimension each). The space of genus 1 curves of degree 3 in 3 space is only 12 dimensional, so most such zero tension curves are not tropicalizations of actual curves. More explicitly, we can embed C into a complete polyhedral complex Σ for which Σ_1 is the fan of projective space. Then $\overline{X}$ will be a degree 3 curve in $\mathbb{P}^3$. Any degree 3, genus 1 curve in $\mathbb{P}^3$ lies in a hyperplane. But, if P, Q and R do not lie on a tropical line, then C is not continued in the tropicalization of any hyperplane, a contradiction.

A zero tension curve is said to be *super abundant* if it has "too large" a space of deformations and *ordinary* if the space of deformations has "the right" dimension which, we will see in Section 5.1 is $E-n(g-3)$ where E is the number of bounded edges of Γ. Mikhalkin conjectured that a regular zero tension curve is

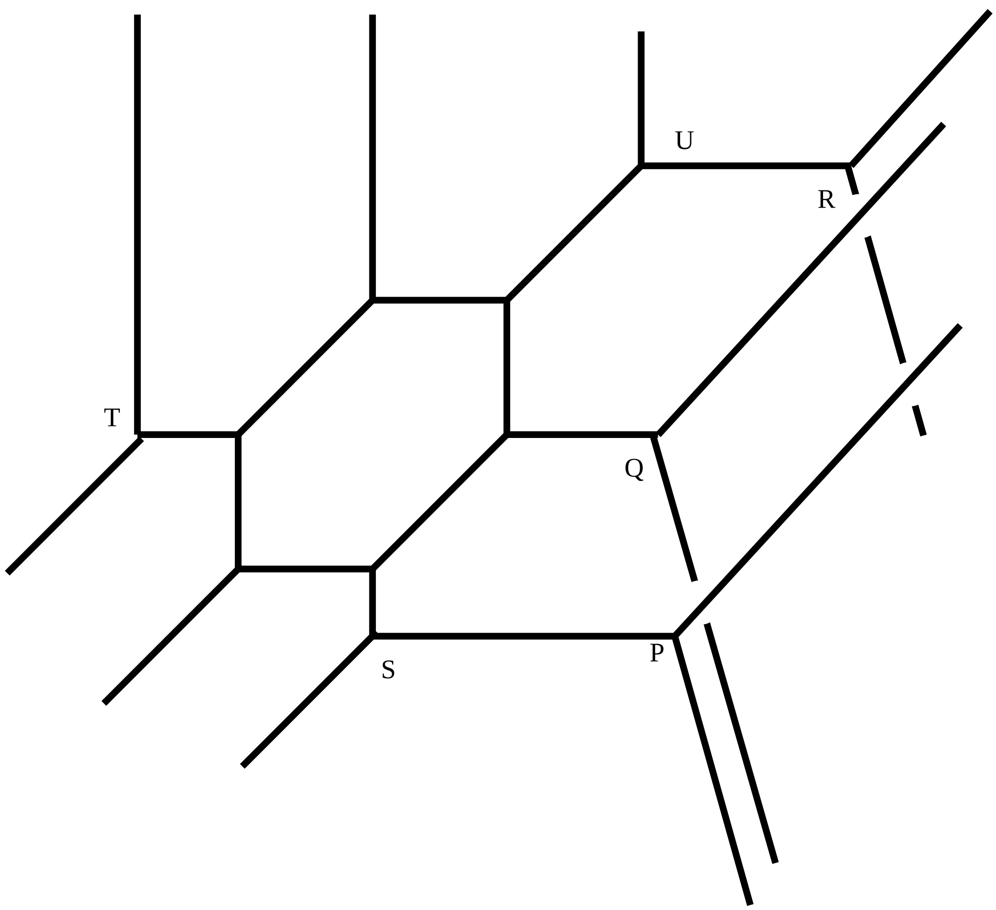

Figure 5.1: A Genus 1 Zero Tension Curve which is not Tropical

always a tropical curve. In this section, we will prove that conjecture completely when κ has characteristic zero, we will need some hypotheses on the characteristic of κ when our curve has genus greater than 2.

Theorem 5.0.4. *If κ has characteristic zero, then every ordinary zero-tension curve is the tropicalization of a curve of corresponding genus and degree. Every ordinary zero-tension curve of genus* 0 *or* 1 *is the tropicalization of a curve of corresponding genus and degree. If Γ is an ordinary zero tension curve and κ has characteristic p then Γ is the tropicalization of a curve of corresponding genus and degree assuming that the matrix* $\mathrm{Slope}(\Gamma, \iota, wt)$ *defined in Section 5.1 has full rank modulo p.*

These characteristic issues are frustrating. I do not have any examples where an ordinary zero tension curve in characteristic p is not achievable as the tropicalization of a curve and very much suspect that none exist. The matrix $\mathrm{Slope}(\Gamma, \iota, wt)$ referred to in the above theorem is a matrix of integers which has full rank if and only if the zero tension curve is ordinary. All of the dependency on characteristic is bound up in Lemma 5.7.1, and we will discuss it more fully at that point.

Mikhalkin has announced a proof of the above conjecture by means of Floer cohomology, which is yet to appear. Our methods will be very different from Mikhlkain's; we use nonarchimedean analysis and uniformizations of curves.

Note that our terminology differs from that of Mikhalkin in [24], who would call something a tropical/parameterized tropical curve if it was a combinatorial object that looked like one of our objects. We reserve the adjective "tropical" to refer to graphs that actually occur as tropicalizations and will refer to combinatorial things that look like tropical curves as "zero-tension".

Finally, we should note that our work bears some interesting resemblances to that of [28] and the sequels to that paper. Morgan and Shalen compactify Teichmuller space, which they think of as the space of two dimensional representations of the free group, by using amoebas and Bergman's notion of a logarithmic limit set. This latter is the same as the tropicalization. Teichmuller space is also closely related to the moduli space of curves, a connection which Morgan and Shalen mention but do not emphasize as they are more concerned with applications to three dimensional hyperbolic manifolds. The main computations in this section deal with studying maps of curves to toric varieties, which Morgan and Shalen do not pursue.

5.1 Combinatorics of Zero Tension Curves

Let Γ be a finite graph. Let Γ^{fin} be the subgraph of Γ consisting of the edges neither of whose endpoints has degree 1; we will call an edge which has an end point of degree 1 an *infinite* edge and an edge which does not have such an endpoint a *finite* edge. Let $w : \mathrm{Edge}(\Gamma) \to \mathbb{Z}_0$ be a function assigning a positive integral weight to each edge. Let $\iota : \Gamma \setminus \{\text{vertices of degree one }\} \to \mathbb{R}^n$ be a continuous map under which each unbounded edge of Γ is taken to a semi-infinite ray, each finite edge is taken to a line segment of finite (nonzero) length and each of these rays and line segments has rational slope. We will use abuse notation by writing $\iota : \Gamma \to \mathbb{R}^n$. For e an edge and v one of its endpoints, let $\tilde{\sigma}_v(e)$ be the primitive lattice vector pointing in the direction of $\iota(e)$ away from $\iota(v)$ and let $\sigma_v(e) = w(e)\tilde{\sigma}_v(e)$. $\sigma_v(e)$ is the important definition, $\tilde{\sigma}_v(e)$ will only occur rarely. Notice that, if e has endpoints v_1 and v_2, then $\sigma_{v_1}(e) = -\sigma_{v_2}(e)$.

We call (Γ, ι, w) a *zero tension curve* if, at every vertex v not of degree 1, we have $\sum_{e \ni v} \sigma_v(e) = 0$. We sometimes will abuse notation by referring to Γ as the zero tension curve, saying, for example, "If a zero tension curve is a tree..." We saw in Section 2.5 that, if $X \subset (K^*)^n$ is a curve then $\operatorname{Trop} X$ naturally has the structure of $\iota(\Gamma)$ for a zero tension curve (Γ, ι, w). The goal of this section is to prove some combinatorial results about zero tension curves. We assume for convenience that Γ is connected.

Suppose that if v is a vertex of Γ of degree 2 with edges e_1 and e_2 coming out of it connected to vertices u_1 and u_2. Then $\sigma_v(e_1) = -\sigma_v(e_2)$. This forces $\iota(e_1)$ and $\iota(e_2)$ to be parallel and point in opposite directions away from $\iota(v)$ (since $w(e_1)$ and $w(e_2) > 0$) so $\tilde{\sigma}_v(e_1) = -\tilde{\sigma}_v(e_2)$ and $w(e_1) = w(e_2)$. Let Γ' be the graph obtained by deleting v, e_1 and e_2 and inserting a single edge e_0 between u_1 and u_2. Define the weight function w' on Γ' by having $w'(e) = w(e)$ for every edge w other than e_0 and $w'(e_0) = w(e_1) = w(e_2)$. Define $\iota' : \Gamma' \to \mathbb{R}^n \cup \{\infty\}$ by having

ι' coincide with ι on the edges other than e_0 and having $\iota'(e_0) = \iota(e_1) \cup \iota(e_2)$. Thus, we have replaced (Γ, ι, w) with a new zero tension curve (Γ', ι', w') which has one fewer vertex. Continuing in this manner, we will always eventually arrive at a zero tension curve with no vertices of degree 2. Since most of our results will be unaffected by this reduction procedure, we will often restrict to the case of curves that have no vertices of degree 2, which we term *non-bivalent*. An especially important case is the case where each vertex is either of degree 3 or 1, which we term *trivalent*.

We define the *degree* of a zero tension curve to be the unordered list of vectors $\sigma_v(e)$ where v runs over the degree 1 vertices. We define the *genus* of a zero tension curve to be the first Betti number of Γ. The main result of this section is

Theorem 5.1.1. *The collection of non-bivalent zero tension curves of given genus and degree in $\mathbb{R}^n$, modulo reparameterization, naturally has the structure of a finite union of rational polyhedral cones (in some $\mathbb{R}^N$).*

We will use the variable x to denote the number of infinite edges of Γ.

Lemma 5.1.2. *Assume Γ is non-bivalent. Let V denote the number of non-degree 1 vertices of Γ and E the number of finite edges. Then $V \leq x + 2(g-1)$ and $E \leq x + 3(g-1)$, with equality in each case if and only if Γ is trivalent.*

Proof. By counting the number of pairs (vertex, edge containing vertex) in two ways, we get $2(E + x) \geq 3V + x$, with equality if and only if Γ is trivalent. Also,

$(E+x)-(V+x)=g-1$ so $E=V+g-1$ and $V=E-g+1$. Substituting these into the previous inequality gives the desired results. $\square$

Define the combinatorial type of a zero tension curve to be underlying abstract graph Γ and the data of the vectors $\sigma_v(e)$.

Proposition 5.1.3. *The set of zero tension curves of given combinatorial type has the structure of the relative interior of a cone. Using the notations V and E for the number of finite edges and finite vertices, this cone (if nonempty) has dimension at least $E-n(g-1)$.*

Proof. This is fairly clear. All that remains to specify a zero tension curve once its combinatorial type is known is to give the coordinate of one vertex (n parameters) and the lengths of the finite edges (E parameters, each required to be positive.) Not every collection of edge lengths is legitimate however: all of the cycles must close up. It is enough to check this for a collection of g cycles forming a homology basis of Γ, we get n equations for each cycle. Thus, we are looking at a slice of an open quadrant of $\mathbb{R}^{n+E}$ by a plane of codimension at most ng, so it is a cone of dimension at least $E-n(g-1)$. $\square$

Remark: When Γ is trivalent, this bound is $x-(n-3)(g-1)$ by the above lemma. If we are considering degree d curves in three space, so $x=4d$, then this quantity is $4d$ which is a standard lower bound for the dimension of the Hilbert scheme of degree d, genus g curves in projective n space. According to

Hartshorne, the first proof of this bound may be in [16], lemma 5, although the result is far older.

We define a combinatorial type to be *ordinary* if this cone has dimension $E - n(g-1)$ and *superabundant* otherwise. Let us note explicitly what it means to be ordinary: choose a basis $\gamma_1, \ldots, \gamma_g$ for $H_1(\Gamma)$. We define a $gn \times E$ matrix $\mathrm{Slope}(\Gamma, \iota, wt)$ whose columns are indexed by the edges of $\Gamma$and whose rows are indexed by ordered pairs (i, γ_j) with $1 \leq i \leq n$, $1 \leq j \leq g$. The $((i, \gamma_j), e)$ entry of $\mathrm{Slope}(\Gamma, \iota, wt)$ is $\sigma_{v_1(e)}(e)_j$ times the (signed) number of times that e appears in γ_i. Clearly, $\mathrm{Slope}(\Gamma, \iota, wt)$ depends only on the combinatorial type of (Γ, ι, wt). (Γ, ι, wt) is ordinary if and only if $\mathrm{Slope}(\Gamma, \iota, wt)$ has full rank.

Remark: Mikhalkin, in [25], seems to suggest that every zero tension curve can be perturbed to a trivalent zero tension curve. This is easy to prove when the curve is ordinary; more specifically, we have:

Proposition 5.1.4. *Let (Γ, ι, w) be a zero tension curve and let v be a vertex of Γ of degree 4 or greater. Form a new graph Γ' by replacing v with two vertices v_1 and v_2, connected to each other by an edge and such that every neighbor of v is connected to either v_1 or v_2 and each of v_1 and v_2 has degree at least 3. It is possible to find a family ι_t of embeddings of Γ' into $\mathbb{R}^n$ as a zero tension curve such that: each edge of $\iota_t(\Gamma')$ other than $\iota((v_1, v_2))$ is parallel to the corresponding edge of $\iota(\Gamma)$, $\iota_t((v_1, v_2))$ shrinks down to $\iota(v)$ as $t \to 0$ and, for every other edge e, $\iota_t(e) \to \iota(e)$ as $t \to 0$.*

Proof. Define a third graph Γ'' by removing the edge (v_1, v_2) from Γ' and adding an unbounded edge at each of v_1 and v_2. There are two cases, based on whether or not Γ'' is connected. If Γ'' is not connected, then we can simply make a family of embeddings of Γ' as in the hypotheses of the theorem by fixing the location of $\iota(v_1)$ at v, keeping every edge other than (v_1, v_2) of constant length and giving the edge (v_1, v_2) length t and the slope forced by the zero tension condition.

We now turn to the case where Γ'' is connected. We consider the family of embeddings of Γ'' of the following combinatorial type: every edge of Γ'' has the same slope as the corresponding edge in Γ except for the two unbounded edges which have the slopes forced by the zero tension condition. Now, Γ'' has exactly as many finite edges as Γ but is of one lesser genus. We can choose our basis of cycles for Γ so that only one of them, say c, when pulled back to Γ', goes along (v_1, v_2). The equations which describe the cones of embeddings of Γ'' of the specified combinatorial type are the same as those for Γ'' plus the equation that says the cycle c closes. We can thus think of the embeddings of Γ as embeddings ι'' of Γ in which the cycle c closes or, equivalently, $\iota''(v_1) = \iota''(v_2)$. Since the combinatorial type of (Γ, ι, w) is ordinary, all of the equations cutting out the cone of embeddings are non-redundant and so the possible directions of $\iota''(v_1) - \iota''(v_2)$ must span $\mathbb{R}^n$. In particular, we can find embeddings ι'' of Γ'' where $\iota''(v_1)$ approaches $\iota''(v_2)$ while $\iota''(v_1) - \iota''(v_2)$ has the slope which the edge (v_1, v_2) of Γ' is forced to have by the zero tension condition. Then deleting the unbounded rays and drawing back in (v_1, v_2) gives a sequences of embeddings of Γ' as desired. ☐

Applying this result over and over again, every ordinary curve can be written as a degeneration of a trivalent curve. In the superabundant case, one often can not choose how to split the edge incident on a vertex of degree ≥ 4 and sometimes can not split one vertex without splitting another vertex elsewhere. Nonetheless, it seems like it may always be true that every zero tension curve can be written as the limit of trivalent zero tension curves.

So the set of nonbivalent zero tension curves of given degree and genus is a union of cones, one for each combinatorial type, and we must simply show that there are finitely many combinatorial types of zero tension curve of each genus and degree. First note that there are only finitely many nonbivalent graphes with x degree one vertices and first Betti number g. (Proof: they are all subgraphs of the complete graph on $(x+2g-2)+x$ vertices.) Thus, it is enough to prove, for a fixed graph Γ and a choice of degree, that there are only finitely many choices for the slopes $\sigma_v(e)$. Explicitly, a choice of degree means specifying the slope $\sigma_v(e)$ for each infinite edge e of Γ. We will prove:

Proposition 5.1.5. *For each i between 1 and n, there are only finitely many possible choices for the integer valued function $(\sigma_v(e))_i$.*

Proof. Fix an orientation of the graph Γ, so that each edge e has a chosen ordering $(v_1(e), v_2(e))$ of its endpoints. Let C_1 be the free abelian group on the edges of

Γ.

Now, suppose we are given an integer valued function $s(e)$ on the finite edges of Γ and a positive real valued function ℓ on the finite edges. Suppose there is a zero tension curve with underlying graph Γ, with $s(e) = (\sigma_{v_1(e)}(e))_i = -(\sigma_{v_2(e)}(e))$ and with $\ell(e)\sigma_{v_1(e)}(e)$ parallel to and the same length as $\iota(e)$. Consider the element $S = \sum_{e \in \Gamma^{\mathrm{fin}}} s(e)e$ of C_1. Let T be the analogous sum over the infinite edges; note that T is fixed by our knowledge of the degree. Letting ∂ denote the map $C_1 \to C_0$ sending $e \mapsto v_1(e) - v_2(e)$, the zero tension condition states that $\partial(S+T) = 0$. Thus, ∂S is determined and we see that the different possible values for S differ by members of $\operatorname{Ker} \partial = H_1(\Gamma)$.

On the other hand, the assumption that the cycles of Γ close up tells us that S is in the orthogonal complement of $H_1(\Gamma)$ with respect to the inner product $<,>_\ell$ under which the finite edges of Γ are orthogonal and have norm $< e, e >_\ell = \ell(e)$. We are therefore reduced to proving the following proposition, with $\mathbb{R}^N = C_1 \otimes \mathbb{R}$, $H = H_1(\Gamma) \otimes R$ and $v_0 = S_0$.

Lemma 5.1.6. *Let H be a sub-vector space of $\mathbb{R}^N$ and let $v_0 \in \mathbb{R}^N$. Let $\mathcal{S}$ denote the set of points $v \in \mathbb{R}^N$ such that (1) $v \in v_0 + H$ and (2) v is orthogonal to H for some inner product which is diagonal and positive definite in the standard basis of $\mathbb{R}^N$. Then $\mathcal{S}$ is bounded. More precisely, form a hyperplane arrangement $\mathcal{A}$ in $v_0 + H$ by intersecting $v_0 + H$ with the coordinate hyperplanes of $\mathbb{R}^N$. Then $\mathcal{S}$ is the interior of the union of the bounded regions of $\mathcal{A}$. In particular, this set* *contains finitely many lattice points.*

Proof. First, suppose that v_1 is not in the interior of the bounded part of $\mathcal{A}$. Then there is a $w \in H \setminus \{0\}$ such that $v_1 + tw$ lies in the same region of $\mathcal{A}$ for all $t > 0$. For every coordinate i for which w_i is nonzero, w_i and $(v_1)_i$ must have the same sign. But then $< v_1, w >_\ell$ is entirely a sum of positive terms and $< v_1, w >_\ell > 0$, contradicting that $v_1 \in H^\perp$.

Conversely, suppose that v_1 is in the interior of of the bounded part of $\mathcal{A}$. Then there is no $w \in H \setminus \{0\}$ such that $v + tw$ lies in the same component of $\mathcal{A}$ for all $t > 0$. Reversing the above, that means that there is no $w \in H$ such that w_i and $(v_1)_i$ have the same sign whenever $w_i \neq 0$. By linear programming duality (see, for example, proposition 6.8 in [51]), this implies that $H^\perp$ contain an element with the same sign pattern as v_1, say v_2. Set $\alpha_i = (v_2)_i/(v_1)_i$ with α_i an arbitrary positive real when $(v_2)_i = (v_1)_i$. Set $\ell'_i = \alpha_i \ell_i$. Then v_1 is orthogonal to H with respect ot $<, >_{\ell'}$. □

□

When $g = 0$, this proof shows that there is a unique combinatorial type given the underlying graph Γ and its degree. When $g = 1$, Γ has a unique loop v_1, $\ldots$, v_r. The above proof shows that $\sigma_v(e)$ is uniquely determined for e not in the loop. Moreover, for each v_i, let d_i be the sum of $\sigma_{v_i}(e)$ summed over $e \ni v_i$ where e is not in the loop, then the possible combinatorial structure are in bijection with the lattice points in the interior of the convex hull of the d_i. When $g \geq 2$, it is easy enough to compute the bounded regions for each coordinate $1 \leq i \leq n$, but I don't know a simple way to figure out when two solutions coming from lattice points from different i can arise from the same ℓ.

5.2 The Bruhat-Tits Tree and Cross-Ratios

In this section we review some standard constructions. A good reference for these results with a view towards the sort of applications we will be making is Chapter 2 of [28]. For simplicity, we will assume that $v : K^* \to \mathbb{R}$ is surjective.

We denote by $\mathrm{BT}(K)$ the set of $\mathcal{R}$-submodules of K^2 which are isomorphic to $\mathcal{R}^2$, modulo K^*-scaling. We write $\overline{M}$ for the equivalence class of a module M. We equip $\mathrm{BT}(K)$ with the metric where $d(\overline{M_1}, \overline{M_2})$ is the infimum of all ϵ such that there exists an α with $M_1 \supseteq t^{\alpha} M_2 \supseteq t^{\epsilon} M_1$; this is easily seen to be independent of the choice of representatives M_1 and M_2. $\mathrm{BT}(K)$ is called the Bruhat-Tits tree of K.

If we made the analogous construction working with the field of power series of integral exponents, we could equip $\mathrm{BT}(K)$ with the structure of the vertices of a tree so that distance was the graph theoretic distance. Instead, $\mathrm{BT}(K)$ is what is called an $\mathbb{R}$-tree (see [28]). The following proposition lists the "tree-like" properties of $\mathrm{BT}(K)$.

Proposition 5.2.1. *If $\overline{M_1}$ and $\overline{M_2} \in \mathrm{BT}(K)$ with $d(\overline{M_1}, \overline{M_2}) = d$ then there is a unique distance preserving map $\phi : [0, d] \to \mathrm{BT}(K)$ with $\phi(0) = \overline{M_1}$ and*

$\phi(d) = \overline{M_2}$. *We will call the image of* ϕ *the path from* $\overline{M_1}$ *to* $\overline{M_2}$ *and denote it by* $[\overline{M_1}, \overline{M_2}]$. *If* $\overline{M_1}, \ldots, \overline{M_n} \subset \mathrm{BT}(K)$ *then* $\cup_{i \neq j}[\overline{M_1}, \overline{M_2}]$ *is a tree.*

It is easy to give an explicit description of ϕ: if $M_1 \supseteq M_2 \supseteq t^d M_1$ then $\phi(e) = \overline{t^e M_1 + M_2}$. Suppose now that $(x_1 : y_1)$ and $(x_2 : y_2)$ are distinct members of $\mathbb{P}^1(K)$. Then we can similarly define a map $\phi : \mathbb{R} \to \mathrm{BT}(K)$ by $\phi(e) = \overline{\mathcal{R}(x_1, y_1) + t^e \mathcal{R}(x_2, y_2)}$. We will call the image of this ϕ the *path from* $(x_1 : y_1)$ *to* $(x_2 : y_2)$ and denote it $[(x_1 : y_1), (x_2 : y_2)]$. Similarly, if $\overline{M} \in \mathrm{BT}(K)$ and $(x_1 : y_1) \in \mathbb{P}^1(K)$, we can define a semi-infinite path from $(x_1 : y_1)$ to M denoted $[(x : y), \overline{M}]$.

If Z is a subset of $\mathrm{BT}(K) \cup \mathbb{P}^1(K)$, we denote by $[Z]$ the subspace $\cup_{z,z' \in Z}[z, z'] \subset \mathrm{BT}(K)$. For simplicity, assume that $|Z| \geq 3$. If Z is finite, $[Z]$ is a tree with a semi-infinite ray for each member of $Z \cap \mathbb{P}^1(K)$. We will say that this ray has its end at the corresponding member of $Z \cap \mathbb{P}^1(K)$. We will abbreviate $[\{x_1, \ldots, z_n\}]$ as $[z_1, \ldots, z_n]$.

The particular case where Z is a four element subset of $\mathbb{P}^1(K)$ is of particular importance. Let $\{w, x, y, z\} \subset \mathbb{P}^1(K) = K \cup \{\infty\}$. We define the *cross ratio* $c(w, x : y, z)$ by

$$c(w, x : y, z) = \frac{(w - y)(x - z)}{(w - z)(x - y)}.$$

Note that $c(w, x : y, z) = c(x, w : z, y) = c(y, z : w, x) = c(z, y : x, w)$.

Proposition 5.2.2. $[w, x, y, z]$ *is a tree with* 4 *semi-infinite rays and either* 1 *or* 2 *internal vertices.*

If $[w, x, y, z]$ *has two internal vertices, let* d *be the length of the internal edge and suppose that the rays ending at* w *and* x *lie on one side of that edge and the rays through* y *and* z *on the other. Then* $v(c(w, x : y, z)) = 0$*,* $v(c(w, y : x, z)) = d$ *and* $v(c(w, y : z, x)) = -d$*. The first statement can be strengthened to say that* $v(c(w, x : y, z) - 1) = d$*. (All other permutations of* $\{w, x, y, z\}$ *can be deduced from this).*

If $[\{w, x, y, z\}]$ *has only one internal vertex then* $v(c(w, x : y, z)) = 0$ *and the same holds for all permutations of* $\{w, x, y, z\}$*.*

This proposition can be remembered as saying "$v(c(w, x : y, z))$ is the signed length of $[w, x] \cap [y, z]$" where the sign tells us whether the two paths run in the same direction or the opposite direction along their intersection.

Proof. Due to the invariance of the definition of $\mathrm{BT}(K)$, the group $\mathrm{GL}_2(K)$ acts on it and, since we only consider submodules of K^2 up to K^*, $\mathrm{GL}_2(K)$ acts through its quotient $\mathrm{PGL}_2(K)$. This action is compatible with the standard action of $\mathrm{PGL}_2(K)$ on $\mathbb{P}^1(K) = K \cup \{\infty\}$. It is well known that c is $\mathrm{PGL}_2(K)$ invariant. So the whole theorem is invariant under $\mathrm{PGL}_2(K)$ and we may use this action to take w, x and y to 0, 1 and ∞. Our hypothesis in the second paragraph is that $[0, 1, \infty, z]$ is a tree with 0 and 1 on one side of a finite edge of length d and z and ∞ on the other. It is easy to check that this is equivalent to requiring that $v(z) = -d < 0$. Then $c(0, 1 : \infty, z) = 1 - 1/z$ which does indeed have valuation 0 and $c(0, 1 : \infty, z) - 1 = -1/z$ does indeed have valuation d. $c(0, \infty : 1, z) = 1/z$

which has valuation d and $c(0, \infty : z, 1) = z$ which has valuation $-d$. In the second paragraph, the assumption that the tree has no finite edge implies that $v(z) = v(z-1) = 0$ and the argument then continues as before. $\square$

5.3 Tropical Genus Zero Curves

The aim of this section is to prove Theorem 5.0.4 in the case where Γ is a tree so we want X to have genus 0. This result will appear in a future publication of Mikhalkin; it also appears with many results on incidence conditions in [30]. Our method of proof is not only far more explicit than these, but it will be a good warm up for the case of higher genus curves. Note that this will generalize the characterization of tropicalizations of lines in Chapter 3. Let (Γ, ι, w) be a zero tension curve with Γ a tree. Put a metric on Γ such that the unbounded edges have infinite length and the length ℓ of a finite edge e is chosen such that $\iota(e)$ is a displacement of $\ell\sigma_v(e)$.

Proposition 5.3.1. *Let T be a metric tree with finitely many vertices such that each degree 1 vertex of T is at the end of an infinite ray. Then there is a $Z \subset \mathbb{P}^1(K)$ such that $[Z]$ is isometric to $[T]$. If we permit some of the degree 1 vertices of T to be at the end of finite length edges then we can still embed T as a subtree of* $\mathrm{BT}(K)$.

Proof. First, we consider the case where each degree 1 vertex of T is at the end of an infinite ray.

Our proof is by induction on the number of finite edges of T. If T has l leaves and no finite edges then T is isometric to $[z_1, \ldots, z_l]$ for $\{z_1, \ldots, z_l\}$ any l elements of K^* with valuation 0 and distinct images in κ^*.

Now, let e be a finite edge of T of length d joining vertices v_1 and v_2. Remove e from T, separating T into two trees T_1 and T_2. Define trees T'_s, where $s = 1, 2$, by adding an unbounded edge to T_s at v_s. By induction, we can find subsets Z_1 and $Z_2 \subset \mathbb{P}^1(K)$ with $[Z_s]$ isometric to T'_s. Let $z_s \in Z_s$ be the element of Z_s at the end of the new ray added to T_s. Without loss of generality, we may assume that $z_1 = 0$ and $z_2 = \infty$. Then the point of T_s corresponding to v_s lies somewhere on $[0, \infty]$. By multiplying Z_1 and Z_2 by elements of K^*, we may assume that these points lie distance d apart with v_1 closer to 0 than v_2 is to 0. Then T is isometric to $[Z_1 \cup Z_2]$.

Finally, if not each degree 1 vertex of T is at the end of an infinite ray then we can embed T into a tree T' which does have this property, find an isometric embedding of T' in $\mathrm{BT}(K)$ by the above and then T will be embedded as a subtree of T'. $\square$

Let $Z \subset \mathbb{P}^1(K)$ be such that $[Z]$ is isometric to Γ. We define multisets $Z_1^+, \ldots, Z_n^+, Z_1^-, \ldots, Z_n^-$ as follows: All of the elements of $Z^{\pm,i}$ lie in Z. Let $z \in Z$ correspond to the end of an infinite ray e of Γ. Suppose that $\sigma_z(e) = (s_1, \ldots, s_n)$. Then $z \in Z^{\pm,i}$ if and only if $\pm s_i < 0$. In this case, $|s_i|$ is the number of times that z occurs in $Z^{\pm,i}$. Define a rational map $\phi : \mathbb{P}^1(K) \to K^n$ by the formula

$$\phi(u) = (\phi_1(u), \ldots, \phi_n(u)) = \left(\frac{\prod_{z \in Z_1^+}(u-z)}{\prod_{z \in Z_1^-}(u-z)}, \ldots, \frac{\prod_{z \in Z_n^+}(u-z)}{\prod_{z \in Z_n^-}(u-z)} \right).$$

Here u is a coordinate on $\mathbb{P}^1(K)$, thought of as $K \cup \{\infty\}$.

Theorem 5.3.2. $\operatorname{Trop} \phi(\mathbb{P}^1(K))$ *is a translation of* $\iota(\Gamma)$.

Clearly, $\phi(\mathbb{P}^1(K))$ is a genus zero curve of the appropriate degree. From now until the end of the proof, we identify $[Z]$ with Γ so that we can write $\iota : [Z] \to \mathbb{R}^n$.

Proof. Let $u \in \mathbb{P}^1(K) \setminus Z$. Then $[Z]$ is a tree and $[Z \cup \{u\}]$ is a tree with one additional end. Let $b(u) \in [Z]$ be the point at which that end is attached. We claim that, up to a translation, $v(\phi(u))$ is $\iota(b(u))$. In other words, if u_1 and u_2 are distinct members of $u \in \mathbb{P}^1(K) \setminus Z$, we must show that for each i between 1 and n we have

$$v(\phi_i(u_1)) - v(\phi_i(u_2)) = \iota(b(u_1))_i - \iota(b(u_2))_i.$$

It is enough to show this in the case where $b(u_1)$ and $b(u_2)$ lie in the same edge e of $[Z]$. We will fix one coordinate i to pay attention to, so i will not appear in our notation. Let $Z_i^+ = \{z_1^+, \ldots, z_r^+\}$ and $Z_i^- = \{z_1^-, \ldots, z_r^-\}$. We may find constants $1 \leq s^+, s^- \leq n$ and order the $z_j^\pm$ such that $z_j^\pm$ is on the $b(u_1)$ side of e for $1 \leq j \leq s^\pm$ and on the $b(u_2)$ side of e for $s^\pm + 1 \leq j \leq r$. Let d be the distance from $b(u_1)$ to $b(u_2)$.

We have

$$\begin{aligned} v(\phi_i(u_1)) - v(\phi_i(u_2)) &= v\left(\frac{\phi_i(u_1)}{\phi_i(u_2)}\right) \\ &= v\left(\frac{\left(\prod_{j=1}^r (u_1 - z_j)^+ / \prod_{j=1}^r (u_1 - z_j)^-\right)}{\left(\prod_{j=1}^r (u_2 - z_j)^+ / \prod_{j=1}^r (u_2 - z_j)^-\right)}\right) \\ &= v\left(\prod_{j=1}^r c(u_1, u_2 : z_j^+, z_j^-)\right) \\ &= \sum_{j=1}^r v(c(u_1, u_2 : z_j^+, z_j^-)) = d(s^+ - s^-). \end{aligned}$$

The last equality is by applying Proposition 5.2.2 to each term.

By the zero tension condition, $s_i(e) = s^+ - s^-$ (recall that the slope of $\iota(e)$ is $(s_1(e), \ldots, s_n(e))$.) So $\iota(b(u_1))_i - \iota(b(u_2))_i$ is also $d(s^+ - s^-)$. $\square$

We pause for two examples.

Example 5.3.3. Consider the tree in $\mathbb{R}^3$ with a finite edge running from $(0,0,0)$ to $(1,1,1)$, infinite edges leaving $(1,1,1)$ in directions $(1,0,0)$ and $(0,1,1)$ and edges departing $(0,0,0)$ in directions $(0,-1,0)$ and $(-1,0,-1)$. Then $[0, t, 1, t^{-1}]$ is isometric to Γ, with 0, t,1 and t^{-1} respectively corresponding to the endpoints of the above infinite rays. We have

$$Z^{+,1} = \{0\} \quad Z^{+,2} = \{t\} \quad Z^{+,3} = \{t\}$$

$$Z^{-,1} = \{t^{-1}\} \quad Z^{-,2} = \{1\} \quad Z^{-,3} = \{t^{-1}\}$$

Thus, the map ϕ is given by

$$u \mapsto \left(\frac{u}{u - t^{-1}}, \frac{u - t}{u}, \frac{u - t}{u - t^{-1}}\right).$$

The image of this map is a genus 0 curve X with $\operatorname{Trop} X$ equal to the given tree.

Example 5.3.4. This time we choose a tree with no internal edges but complcated slopes. Consider the tree T in $\mathbb{R}^3$ with no internal edges and four unbounded rays of slope $(1,2,3)$, $(5,-3,4)$, $(-7,1,-2)$, $(1,0,-5)$. Assuming that κ has characteristic 0, the tree $[1,2,3,4] \subset \mathrm{BT}(K)$ is isometric to T. Our multisets $Z_i^{\pm}$ are

$$Z^{+,1} = \{1,2,2,2,2,2,4\} \quad Z^{+,2} = \{1,1,3\} \quad Z^{+,3} = \{1,1,1,2,2,2,2\}$$

$$Z^{-,1} = \{3,3,3,3,3,3,3\} \quad Z^{-,2} = \{2,2,2\} \quad Z^{-,3} = \{3,3,4,4,4,4,4\}$$

For example, there are 5 occurences of the number 4 in $Z^{-,3}$ because ray number 4 of out tree has slope -5 in the x_3 direction.

Our map ϕ is given by

$$u \mapsto \left(\frac{(u-1)(u-2)^5(u-4)}{(u-3)^7}, \frac{(u-1)^2(u-3)}{(u-2)^3}, \frac{(u-1)^3(u-2)^4}{(u-3)^2(u-4)^5} \right).$$

Once again, the image of ϕ is a genus zero curve whose tropicalization is the given tree.

5.4 Tropical Genus One Curves

Let (Γ, ι, w) be a zero tension curve where Γ is connected with first Betti number 1. This means that Γ has a unique cycle, let $e_1, \ldots, e_r$ be the edges of this cycle and let σ_i be $\sigma(e_i)$.

Our aim in this section is to prove

Theorem 5.4.1. *If* (Γ, ι, w) *is ordinary then there is a genus one curve* $X \in (K^*)^n$ *with* $\operatorname{Trop} X = \iota(\Gamma)$

Let us make the condition of the theorem more explicit:

Proposition 5.4.2. (Γ, ι, w) *is ordinary if and only if the* σ_i *span* $\mathbb{R}^n$*.*

Proof. Consider the equation $\sum \ell_i \sigma_i = 0$ where $\ell_i \in \mathbb{R}$. By definition, $(\Gamma, , \iota, w)$ is ordinary if this equation has solution space of codimension n in $\mathbb{R}^r$. In other words, the kernel of the map taking $\mathbb{R}^r \to \mathbb{R}^n$ *via* $(\ell_1, \ldots, \ell_r) \mapsto \sum \ell_i \sigma_i = 0$ must be $(r-n)$-dimensional or, in other words, the map must be surjective. This precisely says that the σ_i span $\mathbb{R}^n$. □

We use Tate's nonarchimedean uniformizations of elliptic curves. A good reference for this subject is chapter V of [34]. Let $q \in K^*$ with $v(q) > 0$. Tate constructs an elliptic curve E over K with a bijection p from $K^*/q^{\mathbb{Z}}$ to $E(K)$. For z^+, $z^- \in K^*$, define

$$\phi_{z^+,z^-}(u) = \prod_{i=-\infty}^{0} \left(\frac{u/z^+ - q^i}{u/z^- - q^i} \right) \prod_{i=1}^{\infty} \left(\frac{u - q^i z^+}{u - q^i z^-} \right)$$

This product is convergent in the nonarchimedean topology on K for all $u \in K^* \setminus q^{\mathbb{Z}} \cdot (z^+, z^-)$ and $\phi(qu) = (z^+/z^-)\phi(u)$. (Remember that $\lim_{n\to\infty} q^n = 0$ because $v(q) > 0$.) Thus, if $Z^+ = \{z_1^+, \ldots, z_k^+\}$ and $Z^- = \{z_1^-, \ldots, z_k^-\}$ are finite multisubsets of K^* with the same cardinality and $\prod_{i=1}^k (z_i^+/z_i^-) = 1$ then $\phi_{Z^+,Z^-}(u) := \prod_{i=1}^k \phi_{z_i^+,z_i^-}(u)$ is a well defined function on $(K^*/q^{\mathbb{Z}}) \setminus \left(\bigcup_{j=\infty}^{\infty} q^j \cdot \{z_1^+, \ldots, z_k^+, z_1^-, \ldots, z_k^-\} \right)$. Thought of as a function on $E(K)$ with the appropriate points removed, Tate proves that this is a rational function with zeroes at the points $p(z_i^+)$ and poles at $p(z_i^-)$, where p is the projection $K^*/q^{\mathbb{Z}} \to E(K)$.

Let $f_1, \ldots, f_m$ be the set of unbounded rays of Γ and v_i their endpoints. Let $d_j = \sum_{i=1}^m \max(\sigma_{v_i}(f_i)_j, 0) = -\sum_{i=1}^m \min(\sigma_{v_i}(f_i)_j, 0)$. The equality is because, by the zero tension condition, $\sum_{i=1}^m \sigma_{v_i}(f_i) = 0$. Let $Z_1^+, \ldots, Z_m^+$, $Z_1^-, \ldots, Z_m^-$ be multisubsets of $(K^*)^n$ with $|Z_i^+| = |Z_i^-| = d_i$. Let ℓ_i be the length of e_i and let $\ell = \sum \ell_i$. Choose $q \in K^*$ with $v(q) = \ell$.

Let $\tilde{\gimel}$ be the subtree of $\mathrm{BT}(K)$ spanned by $\bigcup_{i=\infty}^{\infty} q^i \left(\bigcup_{j=1}^n \left(Z_j^+ \cup Z_j^- \right) \right)$. Let $\gimel$ be the quotient of $\tilde{\gimel}$ under the translation by $q^{\mathbb{Z}}$. ($\gimel$, pronounced "gimmel", is the Hebrew letter analogous to G and Γ. We need to reserve G for a certain group that will appear in the next section.)

Choose an ordering $(z_i^{\pm,1}, \ldots, z_i^{\pm,d_i})$ of each of the $Z_i^{\pm}$. Consider all of the paths $[z_i^{-,k}, z_i^{+,k}]$ for $1 \leq k \leq d_i$ and take their images down in $\gimel$; call this collection P_i. For an edge e of $\gimel$ with endpoint v, by "the number of signed paths of P_i running along e away from v", we mean to count each path with sign according to whether it runs towards or away from v along e and multiplicity counting how many times it passes through e.

Proposition 5.4.3. *It is possible to choose $Z_i^{\pm}$ and q such that*

1. *There is an isometry $\psi : \gimel \xrightarrow{\sim} \Gamma$.*

2. *If e is an edge of $\gimel$ with end point v then the i^{th} component of $\sigma_v(\psi(e))$ is*

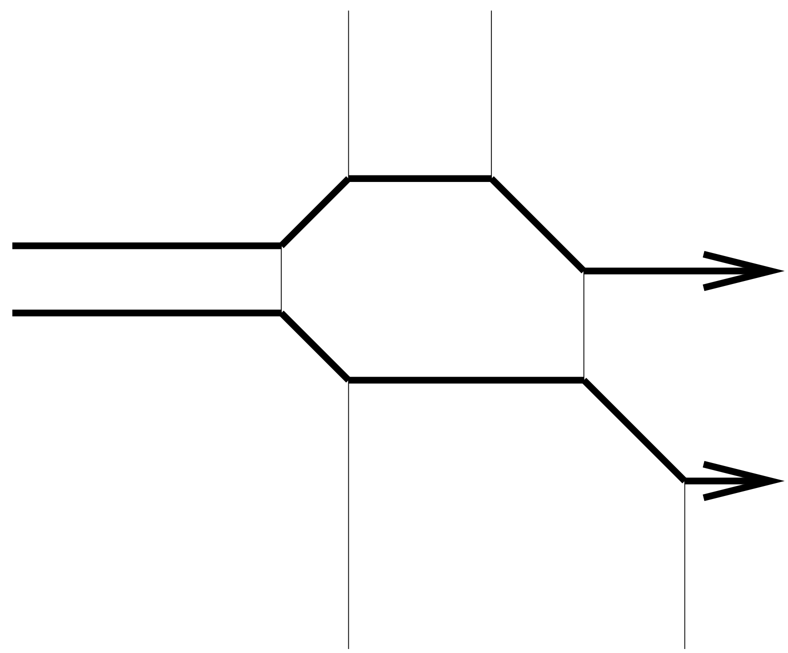

Figure 5.2: A zero tension curve with P_1 marked in bold

the number of signed paths of P_i running along e away from v.

3. *For each $1 \leq i \leq n$, we have $v\left(\prod_{k=1}^{d_i} z_i^{+,k}/z_i^{-,k}\right) = 0$.*

The first condition is better conveyed by a picture. Consider the planar zero tension of genus 1 shown in figure 5.2. (All edges have weight 1.) Let $i = 1$. The pathes P_1 are shown in bold; the first condition says that the number of pathes running through an edge is the first coordinate of its slope.

Proof. Let the vertices of the closed loop of Γ be $v_1, \ldots v_r$ with e_i joining v_i to v_{i+1}. remove e_r from Γ and add two unbounded edges at v_1 and v_r to form a tree T. By Proposition 5.3.1, there is a subtree $[Y']$ of $\mathrm{BT}(K)$ isometric to T, where

Y' is a finite subset of $\mathbb{P}^1(K)$. By an automorphism of $\mathbb{P}^1(K)$, we may arrange that the ends of the new edges added at v_1 and v_r are at 0 and ∞.

Let $Y = Y' \setminus \{0, \infty\}$. We define the multisets $Z^{\pm,i}$ as follows: every element of $Z^{\pm,i}$ is an element of Y. Let $y \in Y$ correspond to the infinite edge e of Γ, then $y \in Z^{\pm,i}$ if and only if $\pm\sigma_y(e)_i < 0$ and the number of times that y occurs is $|\sigma_y(e)_i|$.

Choose q with $v(q) = \ell$. Let $\sigma_{v_r}(e_r) = (s_1, \ldots, s_n)$. Choose $Z_i^{\pm}$ as follows: We take $Z_i^- = Y_i^-$. Z_i^+ is the same as Y_i^+ except that we multiply $y_i^{+,1}$ by q^{s_i}. We claim that the proposition holds for this $Z_i^{\pm}$ and this q. It is clear that Γ and $\gimel$ are isomorphic as graphs, we must check that the corresponding edges have the same lengths to prove claim (1). By constructing $[Y']$ isometric to T, this is automatic for every edge except the ones corresponding to e_r. By our choice to make $v(q) = \ell$, multiplication by q shifts $\tilde{\gimel}$ by ℓ along the path $[0, \infty]$ so the cycle in $\gimel$ has length ℓ. The cycle in $\iota(\Gamma)$ also has length ℓ and all the other edges are the right length, so (1) follows.

For (2), we break into cases depending on the location of the edge e. First, suppose that e is not in the closed loop of $\gimel$. Then removing e from $\gimel$ disconnects e into a tree and a graph with a loop, we may assume that v is in the end that is a tree. Then the number of signed paths of P_i running along e away from v is simply the number of those paths ending in the tree minus the number starting in the tree. By construction, this is the sum of the i components of the directions of the corresponding infinite rays in $\iota(\Gamma)$. By the zero tension condition, these two numbers match.

Next, suppose that $e = e_r$ and $v = v_r$. The only path in P_i crossing e_r is the image of $[y_i^{-,1}, q^{s_i} y_i^{+,1}]$. This path runs along e_r away from v s_i times so (2) holds in this case.

Finally, let $t_v(e)_i$ denote the number of times a path in P_i runs along e away from v. Since every path that enters v must leave it, $\sum_{e \ni v} t_v(e)_i = 0$. In other words, the vectors $(t_v(e)_1, \ldots, t_v(e)_n)$ obey the zero tension condition. The $\sigma_v(e)$, of course, also obey the zero tension condition. But once we have determined $t_v(e)$ on all edges except the chain e_1, $\ldots$, e_{r-1}, the zero tension condition forces these remaining edges to be correct as well.

We now turn to (3). $v(z_i^{+,k}/z_i^{-,k})$ is the length of the overlap between the paths $[z_i^{+,k}, z_i^{-,k}]$ and $[0, \infty]$. Thus $v\left(\prod_{k=1}^{d_i} z_i^{+,k}/z_i^{-,k}\right) = \sum_{k=1}^{d_i} v(z_i^{+,k}/z_i^{-,k})$ is the sum over all k of the length of this overlap. Pushing the paths $[z_i^{+,k}, z_i^{-,k}]$ and $[0, \infty]$ down to $\gimel$, our problem is to sum over the paths of P_i the length of the signed overlap of P_i with the closed loop of $\gimel$. Breaking up our sum over the edges of the closed loop and using property (3), this is $\sum \ell_j \sigma_{v_j}(e_j)_i$ where the subscript i means to take the i component. But $\sum_j \ell_j \sigma_{v_j}(e_j) = 0$ because the loop closes, so we are done. $\square$

Choose q and $Z_i^{\pm}$ as in the above proposition. If we knew that for each i we had $\prod_{z^+ \in Z_i^+} z^+ = \prod_{z^- \in Z_i^-} z^-$ then we would have an embedding of the elliptic curve $K^*/q^{\mathbb{Z}}$ in $(K^*)^n$ via $\phi : u \mapsto (\phi_{Z_1^+, Z_1^-}(u), \ldots, \phi_{Z_n^+, Z_n^-}(u))$.

Proposition 5.4.4. *If we have $\prod_{z^+ \in Z_i^+} z^+ = \prod_{z^- \in Z_i^-} z^-$, define the genus one curve $X \subset (K^*)^n$ by the above embedding. Then* $\operatorname{Trop} X$ *is a translate of* $\iota(\Gamma)$. *(So by translating X by an element of $(K^*)^n$ we can arrange that* $\operatorname{Trop} X = \iota(\Gamma)$.*)*

Proof. Let u_1 and $u_2 \in K^* \setminus \bigcup_{i=\infty}^{\infty} q^i \left(\bigcup_{j=1}^{n} \left(Z_j^+ \cup Z_j^- \right) \right)$. Consider the tree $[\{u_s\} \cup \bigcup_{i=\infty}^{\infty} q^i \left(\bigcup_{j=1}^{n} \left(Z_j^+ \cup Z_j^- \right) \right)] \subset \mathrm{BT}(K)$, where $s = 1, 2$. This tree contains $\tilde{\daleth}$ and one additional unbounded edge, let $b(u_i)$ be the point of $\tilde{\daleth}$ at which the new unbounded edge is attached. We claim that

$$\phi(u_2) - \phi(u_1) = \iota(\psi(\pi(b(u_2)))) - \iota(\psi(\pi(b(u_1))))$$

where π is the projection $\tilde{\daleth} \to \daleth$ and ψ is the identification $\daleth \to \Gamma$. As u_2 varies over $K^* \setminus \bigcup_{i=\infty}^{\infty} q^i \left(\bigcup_{j=1}^{n} \left(Z_i^+ \cup Z_i^- \right) \right)$, the point $b(u_2)$ varies over $\daleth$ so $\phi(u_2)$ sweeps out a translate of $\iota(\psi(\daleth)) = \iota(\Gamma)$, as desired.

We now check the claim. It is enough to consider the case where $b(u_1)$ and $b(u_2)$ lie on the same edge e of $\daleth$, say at distance d apart. We can check the claim coordinate by coordinate; we focus on the i coordinate, so we must compute $v(\phi_{Z^{+,i}, Z^{-,i}}(u_2)) - v(\phi_{Z^{+,i}, Z^{-,i}}(u_1))$. We have

$$\frac{\phi_{Z^{+,i}, Z^{-,i}}(u_2)}{\phi_{Z^{+,i}, Z^{-,i}}(u_1)} =$$

$$\prod_{k=1}^{d_i} \left(\prod_{j=-\infty}^{0} \frac{(u_2/z_k^{+,i} - q^j)(u_1/z_k^{-,i} - q^j)}{(u_2/z_k^{-,i} - q^j)(u_1/z_k^{+,i} - q^j)} \prod_{j=1}^{\infty} \frac{(u_2 - q^j z_k^{+,i})(u_1 - q^j z_k^{-,i})}{(u_2 - q^j z_k^{-,i})(u_1 - q^j z_k^{+,i})} \right)$$

$$= \prod_{k=1}^{d_i} \prod_{j=-\infty}^{\infty} \frac{(u_2 - q^j z_k^{+,i})(u_1 - q^j z_k^{-,i})}{(u_2 - q^j z_k^{-,i})(u_1 - q^j z_k^{+,i})}$$
$$= \prod_{k=1}^{d_i} \prod_{j=-\infty}^{\infty} c(u_2, u_1 : q^j z_k^{i,+}, q^j z_k^{i,-}).$$

(In the nonarchimedean topology all convergent products and sums are absolutely convergent, so we may rearrange freely.) So

$$v(\phi_{Z_i^+, Z^-, i}(u_2)) - v(\phi_{Z_i^+, Z^-, i}(u_1)) = \sum_{k=1}^{d_i} \sum_{j=-\infty}^{\infty} v(c(u_2, u_1 : q^j z_k^{i,+}, q^j z_k^{i,-})).$$

By Proposition 5.2.2, each summand is the signed length of the overlap of $[u_2, u_2]$ and $[q^j z_k^{i,+}, q^j z_k^{i,-}]$. Now, $[q^j z_k^{i,+}, q^j z_k^{i,-}]$ is contained in $\gimel$ and $[u_1, u_2]$ meets $\gimel$ only in a length d segment of edge e. So each term is either d, 0 or $-d$ depending on whether or not $[q^j z_k^{i,+}, q^j z_k^{i,-}]$ passes through e and in which direction. Now $[q^j z_k^{i,+}, q^j z_k^{i,-}]$ passes through e if and only if $[z_k^{i,+}, z_k^{i,-}]$ passes through $q^{-j} e$. The number of j for which the latter occurs (counted with sign) is the same as the (signed) intersection of $\pi(e)$ with $\pi([z_k^{i,+}, z_k^{i,-}])$. Summing over k, this is precisely the number of paths in P_i containing e, which is, by assumption, $\sigma_v(\pi(e))_i$ where v is the endpoint of $\pi(e)$ closer to $\pi(b(u_1))$. Then we have $v(\phi_{Z_i^+, Z^-, i}(u_2)) - v(\phi_{Z_i^+, Z^-, i}(u_1)) = d\sigma_v(\pi(e))_i$, which is indeed the i component of $\iota(\psi(\pi(b(u_2)))) - \iota(\psi(\pi(b(u_1))))$. $\square$

We now get to the heart of the proof, and to the stage which will be far more difficult when we consider curves of higher genus.

Proposition 5.4.5. *Suppose that* (Γ, ι, w) *is ordinary. It is possible to choose*

$Z_i^{\pm}$ and q so that $\daleth$ is isometric to Γ and $\prod_{z^+ \in Z_i^+} z^+ = \prod_{z^- \in Z_i^-} z^-$ for every $1 \leq i \leq n$.

Proof. We start with a provisional choice of $Z_i^{\pm}$ and q as in Proposition 5.4.3. For $1 \leq j \leq r$, let $Z_i^{\pm}(j)$ be the submultiset of $Z_i^{\pm}$ consisting of those z for which the corresponding degree 1 vertex of $\daleth$ connects to the closed loop at v_j. Let u_1, $\dots$, $u_r \in K^*$ be such that $v(u_j) = 0$ but otherwise arbitrary. We will make a modified choice of $Z_i^{\pm}$ by multiplying each member of $Z_i^{\pm}(j)$ by u_j. This has the effect of changing $\prod_{k=1}^{d_i} z_i^{k,+}/z_i^{k,i}$ by $\prod_{j=1}^{r} u_j^{|Z_i^+(j)|-|Z_i^-(j)|}$.

The zero tension principle tells us that we have $|Z_i^+(j)| - |Z_i^-(j)| + \sigma_{v_j}(e_j)_i + \sigma(v_j)(e_{j-1})_i = 0$. In other words, $\sigma(v_{j-1})(e_{j-1})_i - \sigma_{v_j}(e_j)_i = |Z_i^+(j)| - |Z_i^-(j)|$. We need to show that, by taking u_1, $\dots$, $u_r \in v(K^*)$ with valuation 0 but otherwise arbitrary, we can make the n products $\prod_{j=1}^{r} u_j^{\sigma(v_{j-1})(e_{j-1})_i - \sigma_{v_j}(e_j)_i}$ take on any n values with valuation 0. Since $\{x \in K^* : v(x) = 0\}$ is a divisible group (K is algebraically closed) it is enough to show that the matrix whose (i, j) entry is $(\sigma(v_{j-1})(e_{j-1})_i - \sigma_{v_j}(e_j)_i)$ has rank n over $\mathbb{R}$. In other words, we must show that the vectors $\sigma_{v_{j-1}}(e_{j-1}) - \sigma_{v_j}(e_j)$ span $\mathbb{R}^n$.

Suppose that the $\sigma_{v_{j-1}}(e_{j-1}) - \sigma_{v_j}(e_j)$ do not span $\mathbb{R}^n$. Then there is a nonzero linear function λ on $\mathbb{R}^n$ which takes the value 0 on each $\sigma_{v_{j-1}}(e_{j-1}) - \sigma_{v_j}(e_j)$. So $\lambda(\sigma_{v_i}(e_i))$ is a constant λ_0 independent of i. There are two cases. If $\lambda_0 \neq 0$ then we may assume $\lambda_0 > 0$ and we have $\lambda(\sum_{i=1}^{r} \ell_i \sigma_{v_i}(e_i)) = \lambda \sum_{i=1}^{r} \ell_i > 0$ because each ℓ_i is positive. But $\sum_{i=1}^{r} \ell_i \sigma_{v_i}(e_i) = 0$, as the loop of Γ closes.

On the other hand, if $\lambda_0 = 0$ then the $\sigma_{v_i}(e_i)$ do not span $\mathbb{R}^n$, contradicting, according to Proposition 5.4.2, the assumption that Γ is ordinary. $\square$

Now, choose $Z_i^{\pm}$ and q in the manner that Proposition 5.4.5 guarantees we can. Then the hypothesis of Proposition 5.4.4 is met and we can define a genus one curve $X \subset (K^*)^n$ by the embedding $\phi : K^*/q^{\mathbb{Z}} \hookrightarrow (K^*)^n$. Proposition 5.4.4 tells us that the tropicalization of this curve will be $\iota(\Gamma)$. This proves Theorem 5.4.1.

5.5 Mumford Curves

To prove Theorem 5.0.4 for $g \geq 2$, we will need a way of uniformizing higher genus curves similar to the products of cross ratios used in the preceding proofs. Our tool for this purpose will be the theory of Mumford curves. We will present an extremely brief description of this theory. For more background, see [29], [26] and [19].

Let $\gamma \in \mathrm{PGL}_2(K)$. We say γ is *hyperbolic* if (a representative of) γ is diagonalizable with eigenvalues λ_1 and and λ_2 such that $v(\lambda_1) \neq v(\lambda_2)$. We choose our notation so that $v(\lambda_1) > v(\lambda_2)$ and we will denote by $V_1(\gamma)$ and $V_2(\gamma)$ the images in $\mathbb{P}^1(K)$ of the corresponding eigenvectors of γ. The $V_i(\gamma)$ are fixed points for the action of γ on $\mathbb{P}^1(K)$, every point of $\mathbb{P}^1(K) \setminus \{V_1(\gamma), V_2(\gamma)\}$ tends towards $V_2(\gamma)$ under forward iteration of γ and towards $V_1(\gamma)$ under backwards iteration.

We say that a subgroup $G \subset \mathrm{PGL}_2(K)$ is *Schottky* if G is finitely generated and γ is hyperbolic for every $\gamma \in G \setminus \{\left(\begin{smallmatrix} 1 & 0 \\ 0 & 1 \end{smallmatrix}\right)\}$. Fix a Schottky group G. Let $\Sigma = \overline{\bigcup_{\gamma \in G \setminus \{\left(\begin{smallmatrix} 1 & 0 \\ 0 & 1 \end{smallmatrix}\right)\}} \{V_1(\gamma), V_2(\gamma)\}}$ where $\overline{S}$ denotes the closure of S in the nonarchimedean topology on $\mathbb{P}^1(K)$.

Proposition 5.5.1. $[\Sigma]$ *is a tree.* G *acts freely on* $\mathrm{BT}(K)$ *preserving* Σ*. As a corollary,* G *is free.*

Let g denote the number of generators of the free group G. Let $\Omega = \mathbb{P}^1(K) \setminus \Sigma$. Conceptually, the theory of Mumford curves allows us to define the structure of a smooth curve of genus g on Ω/G.

More precisely, let Z^+ and Z^- be finite multisubsets of Ω with $|Z^+| = |Z^-| = r$. Write the members of $Z^\pm$ as $\{z_1^\pm, \ldots, z_r^\pm\}$. Set $Z = Z^+ \cup Z^-$.

We now define a function $w_{Z^+,Z^-} : \Omega \to K \cup \{\infty\}$ by the infinite product

$$w_{Z^+,Z^-}(u) = \prod_{\gamma \in G} \prod_{j=1}^{j} \frac{u - \gamma(z_j^+)}{u - \gamma(z_j^-)}.$$

This product is convergent in the nonarchimedean topology on K^* and does not depend on the order of the product.

Proposition 5.5.2. *For any* Z^+ *and* Z^- *as above and any* $\gamma \in G$ *there is a unique constant* $\mu_{Z^+,Z^-}(\gamma) \in K^*$ *such that* $w_{Z^+,Z^-}(\gamma u) = \mu_{Z^+,Z^-}(\gamma) w_{Z^+,Z^-}(u)$ *for all* $u \in \Omega$*.*

Theorem 5.5.3 (Mumford). *Let* $\mathcal{K}$ *denote the* K*-algebra of functions on* Ω *generated by the functions* w_{Z^+,Z^-} *coming form those pairs* (Z^+, Z^-) *for which*

μ_{Z^+,Z^-} *is identically* 1*. Then* $\mathcal{K}$ *is a field and there is a smooth genus* g *curve* X *over* K *such that* $\mathcal{K}$ *is the function field of* X*. The points of* X *over* K *are in bijection with* Ω/G*. Every element of* $\mathcal{K}$ *is of the form* aw_{Z^+,Z^-} *for* $a \in K$ *and* μ_{Z^+,Z^-} *identically* 1*.*

Consider $\tilde{\gimel}(G,Z) = [\Sigma \cup \bigcup_{\gamma \in G} \gamma(Z)]$; $\tilde{\gimel}(G,Z)$ is a locally finite tree. G acts freely on $\tilde{\gimel}(G,Z)$; let $\gimel(G,Z)$ denote the quotient graph $\tilde{\gimel}(G,Z)/G$. We denote by π the projection $\tilde{\gimel}(G,Z) \to \gimel(G,Z)$. $\gimel(G,Z)$ is a graph with finitely many vertices and edges. After fixing a base point, we have a natural isomorphism $\pi_1(\gimel(G,Z)) \cong G$. As a corollary, $\gimel(G,Z)$ has first Betti number g. The leaves of $\gimel$ are naturally labeled by the elements of Z.

Let Edge$(\tilde{\gimel}(G,Z))$ and Edge$(\gimel(G,Z))$ denote the $\mathbb{R}$-vector spaces generated by the directed edges of $\tilde{\gimel}(G,Z)$ and $\gimel(G,Z)$ respectively. Here reversing the direction of e is considered to negate it. Define symmetric inner products on these spaces by defining $< e, e' >$ to be 0 if e and e' are distinct edges or if e is of infinite length and $< e, e >= \ell(e)$ otherwise. We will abuse notation by writing a path in one of these graphs to mean the sum of its edges.

Proposition 5.5.4. *Let* u *be any vertex of* $\tilde{\gimel}(G,Z)$*. We have*

$$v(\mu_{Z^+,Z^-}(\gamma)) =< [u, \gamma(u)], \sum_{i=1}^{r} [z_i^+, z_i^-] > .$$

In other words, let L *be any loop in* $\Gamma(G,Z)$ *which realizes the conjugacy class of*

γ in $\pi_1(\gimel(G,Z))$. Then

$$v(\mu_{Z^+,Z^-}(\gamma)) =< L, \sum_{i=1}^{r} \pi([z_i^+, z_i^-]) >$$

where π is the projection $\tilde{\gimel}(G,Z) \to$ $gimel(G,Z)$.

See [26] for a proof. (This is also not that hard to prove by hand, using Proposition 5.2.2.)

5.6 Constructing a Tropical Curve: First Steps

Now let (Γ, ι, w) be a zero tension curve in $\mathbb{R}^n$ and let g be the first Betti number of Γ. We will attempt to construct a genus g curve X over K and n meromorphic functions $\phi_1, \ldots, \phi_n : X \to K \cup \{\infty\}$ such that $\operatorname{Trop} \phi(X) = \iota(\Gamma)$. (Of course, we may fail, as we have not assumed yet that (Γ, ι, w) is ordinary.)

Let Y be the set of endpoints of infinite edges of Γ. We define multisets $Y^{\pm,i}$ whose elements are all elements of Y as before. That is: let $y \in Y$ with y at the end of edge $e \in \Gamma$. Then occurs in $Y^{\pm,i}$ if $\pm\sigma_y(e)_i < 0$ and occurs in that case with multiplicity $|\sigma_y(e)_i|$. By the zero tension condition, $|Y^{+,i}| = |Y^{-,i}|$, call this number r_i.

Proposition 5.6.1. *Fix a coordinate i, with $1 \leq i \leq n$. There are paths P_1, $\ldots P_{r_i} \subset \Gamma$ such that P_j is a path from $y_j^{-,i}$ to $y_j^{+,i}$ and $< L, \sum_{j=1}^{r} P_j >= 0$ for every closed loop $L \subset \Gamma$. Moreover, we can assume that no P_j uses any edge twice.*

Proof. Choose an orientation of Γ; let $v_1(e)$ denote the first endpoint of the edge e. First, note that

$$< L, \sum_e \sigma_{v_1(e)}(e)_i e >= \sum_{e \in L} \ell(e) \sigma_{v_1(e)}(e)_i.$$

This second sum is the net displacement in the i coordinate as we travel around L, which is 0. So it is enough to show that we can find P_j with the specified endpoints such that $\sum_{j=1}^r P_j = \sum_e \sigma_{v_1(e)}(e) e_i$.

No assume that our orientation is such that each edge of Γ is directed from the end with lesser i-coordinate to the end with greater i-coordinate, breaking ties in an acyclic but otherwise arbitrary manner. So $\sigma_{v_1(e)}(e)_i \geq 0$ for every edge e. Give each edge e of this directed graph weight $s_{v_1(e)}(e)_i$. By the zero tension condition, this is a flow, with sources of weight 1 at the $y_j^{-,i}$'s and sinks of weight 1 at the $y_j^{+,i}$'s. Any flow can be written as a sum of paths from sources to sinks. As each P_j respects the direction of $\tilde{\Gamma}$ and $\tilde{\Gamma}$ is acyclic, no P_j can use any edge twice. ◻

Let $\tilde{\Gamma}$ be the universal cover of Γ and let $\tilde{P}_j$ be paths lifting the P_j. Let $\tilde{y}_j^{\pm,i}$ be the appropriate endpoints of the $\tilde{P}_j$. We will show in the next section:

Proposition 5.6.2. *There is a Schottky group G and multisubsets $Z^{\pm,1}$, ..., $Z^{\pm,n}$ of Ω such that, writing $Z = \bigcup_{i=1}^n Z^{\pm,i}$, the tree $\tilde{\beth}(G,Z)$ is isometric to $\tilde{\Gamma}$, $\beth(G,Z)$ is isometric to Γ and the ends of $\tilde{\beth}(G,Z)$ ending at the $\tilde{Z}^{\pm,i}$ correspond to the ends of $\tilde{\Gamma}$ ending at the $\tilde{Y}^{\pm,i}$.*

Proposition 5.6.3. *With G and $Z^{\pm,i}$ as in the above proposition, we have $v(\mu_{Z^{+,i},Z^{-,i}}(\gamma)) = 0$ for every i between 1 and n and every $\gamma \in G$.*

Proof. This is a trivial consequence of propositions 5.5.4 and 5.6.1. □

If we were lucky, we would have $\mu_{Z^{+,i},Z^{-,i}}(\gamma) = 1$, a much stronger claim than merely having valuation zero. Achieving this will be the hardest part of our proof of Theorem 5.0.4. Once we have achieved this, the following result tells us we will be done.

Proposition 5.6.4. *Let G and $Z^{\pm,i}$ be as above. Suppose in addition that, for each i between 1 and n, we have $\mu_{Z^{+,i},Z^{-,i}}(G) \equiv 1$. Map $\Omega/G \to K^n$ by*

$$u \mapsto (w_{Z^{+,1},Z^{-,1}}(u), \dots, w_{Z^{+,n},Z^{-,n}}(u)).$$

Call this map ϕ and its components ϕ_i. Then the image of Ω/G is a genus g curve with tropicalization a translation of $\iota(\Gamma)$.

We abbreviate $\tilde{\mathfrak{T}}(G, Z)$ to $\tilde{\mathfrak{T}}$ and let $\psi : \tilde{\mathfrak{T}} \to \tilde{\Gamma}$ be the isometry.

Proof. Let $u \in \Omega \setminus \bigcup_{\gamma \in G} \gamma Z$. Then $[\tilde{\mathfrak{T}} \cup u]$ is a tree which differs from $\tilde{\mathfrak{T}}$ by adding one new end; let $b(u) \in \tilde{\Gamma}$ be the base of this end. We will show that, for u_1 and $u_2 \in \Omega \setminus \bigcup_{\gamma \in G} \gamma Z$,

$$v(\phi_i(u_1)) - v(\phi_i(u_2)) = \iota(\pi(\psi(b(u_1))))_i - \iota(\pi(\psi(b(u_2))))_i$$

where π is the projection $\tilde{\Gamma} \to \Gamma$.

As in the proof in the genus one case, we have

$$v(\phi_i(u_1)) - v(\phi_i(u_2)) = \sum_{\gamma \in G} \sum_{j=1}^{r} v(c(u_1, u_2 : \gamma z_j^+, \gamma z_j^-)).$$

Also as in the genus one proof, we may assume that $b(u_1)$ and $b(u_2)$ lie on a single edge $e \in \tilde{\Gamma}$.

Since each P_j uses no edge twice, the paths $\gamma \tilde{P}_j$ are disjoint as γ ranges over G and, in particular, at most one passes through e. If $\gamma \tilde{P}_j$ does not pass through e then $v(c(u_1, u_2 : \gamma z_j^+, \gamma z_j^-)) = 0$ and may be dropped from the sum. There will be some $\gamma \tilde{P}_j$ passing through e if and only if P_j passes through e. Since we assumed that $\sum_{j=1}^{r} P_j = \sum s_i(e)e$, the number of such terms is $s_i(e)$. Letting d denote the distance between $b(u_1)$ and $b(u_2)$, the left hand side of the above equation is $ds_i(e)$. This is also the right hand side.

We have now shown that

$$v(\phi_i(u_1)) - v(\phi_i(u_2)) = \iota(\pi(\psi(b(u_1))))_i - \iota(\pi(\psi(b(u_2))))_i$$

and thus that $v(\phi(\Omega))$ is a translation of $\iota(\Gamma)$. As all of the $w_{Z^{+,i}, Z^{-,i}}(u)$ lie in $\mathcal{K}$, the image of ϕ will be a quotient Y of the curve X from Mumford's Theorem. Thus, it will be a tropicalization of a curve of genus $\leq g$. Since $\operatorname{Trop} Y = \iota(\Gamma)$, we also have that Y has genus at least the first Betti number and hence has genus $\geq g$. So Y does have genus g as well and, since $g \geq 2$, this means $Y = X$. $\square$

We thus see that we will be done if we show that we can choose G and $Z_i^{\pm}$ as in Proposition 5.6.2 and such that $\mu_{Z_i^+, Z_i^-}(G) = 1$. We will tackle this in the next two sections.

5.7 Two Lemmas

We pause for two lengthy but important lemmas on the behavior of various aspects of the previous section under perturbation. We write U for the multiplicative group $\{x \in K^* : v(x) = 0\}$.

Lemma 5.7.1. *Let A by an $F \times E$ integer matrix of rank F and let c be a positive real number. If κ has characteristic p, assume that A also has rank F after reduction modulo p. Let $\rho : U^E \times U^F \to U^F$ be an action of U^E on U^F given by convergent power series such that*

$$(\rho(e_1, \ldots, e_E)(f_1, \ldots, f_F))_j = \left(\prod_{i=1}^{E} e_i^{A_{ji}} \right) \left(1 + \sum_{i=1}^{E} O(t^c(e_i - 1)) \right) f_j$$

Then ρ acts transitively.

Here we write $f = O(g)$ for f and g functions to K to express $v(f) \geq v(g)$ and extend the notation in the manner familiar from the archimedean $O(\cdot)$.

Proof. First, note that saying A has rank F is equivalent to saying that the action of $(\kappa^*)^E$ on $(\kappa^*)^F$ by

$$((e_1, \ldots, e_E) \cdot (f_1, \ldots, f_F))_j = \left(\prod_{i=1}^{E} e_i^{A_{ji}} \right) f_j \qquad 1 \leq j \leq F$$

is transitive. We see that every orbit of U^E on U^F contains a point $(f_1, \ldots, f_F)$ with $v(f_j - 1) > 0$ for each j. Let $f^0 = (f_1, \ldots, f_F)$ be such a point and let

$m = \min_{1 \le j \le F} v(f_j - 1)$. We will construct, by induction on k, a sequence $e^k = (e_1^k, \dots, e_E^k)$ such that, writing $\rho(e^k)(f^{k-1}) = f^k$, we have $v(f_k^i - 1) \ge (k+1)\min(m,c)$ and $v(e_i^k - 1) \ge k\min(m,c)$. Thus $\prod_{k=1}^{\infty} e_i^k$ will converge and we will have $\rho(\prod_{k=1}^{\infty} e^k)(f^0) = (1, \dots, 1)$.

Suppose that $e^1, \dots, e^{k-1}$ have been constructed and thus we know what f^{k-1} is. Since A has rank F, there is an $E \times F$ matrix of rational numbers such that $AB = \mathrm{Id}$. Take

$$e_i^k = 1 - \sum_{j=1}^{F} B_{ij}(f_j^{k-1} - 1).$$

This formula makes sense because the B_{ij} do not have p's in their denominators. Moreover, since $v(B_{ij}) \ge 0$, we have that $v(e_i^k - 1) \ge \min_j v(f_j^{k-1} - 1)$. By our inductive assumption we indeed have $v(e_i^k - 1) \ge k\min(m,c)$. Then we have

$$\begin{aligned} f_l^k &= \Big(\rho(e^k) f^{k-1}\Big)_l \\ &= \left(\prod_{i=1}^{E}\Big(1 - \sum_{j=1}^{F} B_{ij}(f_j^{k-1} - 1)\Big)^{A_{li}}\right)\left(1 + \sum_{i=1}^{E} O(t^c(e_i^k - 1))\right) f_l^{k-1} \\ &= \left(1 - \sum_{i=1}^{E} A_{li} \sum_{j=1}^{F} B_{ij}(f_j^{k-1} - 1) + \sum_{j=1}^{F} O(f_j^{k-1} - 1)^2\right) \\ &\qquad\qquad \times \left(1 + \sum_{i=1}^{E} O(t^c(e_j^k - 1))\right) f_l^{k-1} \end{aligned}$$

Interchange the order of the double sum to get $\sum_{j=1}^{F}\sum_{i=1}^{E} A_{li}B_{ij}(f_j^{k-1} - 1)$. By our assumption that $AB = \mathrm{Id}$, we have $\sum_{i=1}^{E} A_{li}B_{ij} = \delta_{lj}$ so the sum is just $(f_l^{k-1} - 1)$. Plugging this in, we have

$$\begin{aligned}
f_l^k &= \left(1-(f_l^{k-1}-1)+\sum_{j=1}^{F} O(f_j^{k-1}-1)^2\right)\left(1+\sum_{i=1}^{E} O(t^c(e_i^k-1))\right) f_l^{k-1} \\
&= \left(1-(f_l^{k-1}-1)+\sum_{j=1}^{F} O(f_j^{k-1}-1)^2\right)\left(1+(f_l^{k-1}-1)\right) \\
&\qquad\qquad\qquad\qquad\qquad\qquad \times\left(1+\sum_{i=1}^{E} O(t^c(e_i^k-1))\right) \\
&= 1+\sum_{j=1}^{F} O(f_j^{k-1}-1)^2+\sum_{i=1}^{E} O(t^c(e_i^k-1))
\end{aligned}$$

We have

$$v(O(f_j^{k-1}-1)^2) \geq 2v(f_j^{k-1}-1) \geq 2k\min(m,c) \geq (k+1)\min(m,c).$$

Here the second inequality is by our inductive assumption. Similarly,

$$v(O(t^c(e_i^k-1))) \geq c+v(e_i^k-1) \geq c+k\min(m,c) \geq (k+1\min(m,c).$$

So f^k obeys the inductively required inequalities.

$\square$

It might seem at first that this proof would not work at all if A did not have full rank modulo p. In fact, the situation is not that bad – one can replace the formula $e_i^k = 1-\sum_{j=1}^{F} B_{ij}(f_j^{k-1}-1)$ by $e_i^k = \prod_{j=1}^{F}(f_j^{k-1})^{B_{ij}}$ and at least have a well defined expression. The trouble is estimating the size of e_i^k-1. For example, if K is the algebraic closure of $\kappa((t))$ where κ has characteristic p, then $v((1+t)^{1/p}-1)=1/p$, which is substantially smaller than $v(t)=1$. That said, even in the generality of this lemma (as opposed to the specific action to which we will apply it in the next section), I don't know any example where the characteristic hypotheses are necessary.

Lemma 5.7.2. *Let $u \in U$. Define an permutation τ_u of $\mathbb{P}^1(K)$ by $\tau_u(z) = uz$ if $v(z) > 0$ and $\tau_u(z) = z$ otherwise. Let w, x, y and z be four distinct points of $\mathbb{P}^1(K) \setminus U$.*

If w, x, y and z all lie in $v^{-1}(\mathbb{R}_{>0})$ then

$$c(w, x : y, z) = c(\tau_u(w), \tau_u(x) : \tau_u(y), \tau_u(z)).$$

The same holds if w, x, y and z all lie in $v^{-1}(\mathbb{R}_{<0})$.

Suppose that $w \in v^{-1}(\mathbb{R}_{>0})$ and x, y and $z \in v^{-1}(\mathbb{R}_{<0})$. Let c denote the minimum of $-v(y)$ and $-v(z)$. Then

$$\frac{c(\tau_u(w), \tau_u(x) : \tau_u(y), \tau_u(z))}{c(w, x : y, z)} = 1 + O(t^c(u-1)).$$

The same holds if $w \in v^{-1}(\mathbb{R}_{<0})$ and x, y and $z \in v^{-1}(\mathbb{R}_{>0})$ where c is the minimum of $v(y)$ and $v(z)$.

Suppose that w and x are in $v^{-1}(\mathbb{R}_{>0})$ and y and z are in $v^{-1}(\mathbb{R}_{<0})$. Set $c = \min(-v(y), -v(z)) + \min(v(w), v(x))$. Then

$$\frac{c(\tau_u(w), \tau_u(x) : \tau_u(y), \tau_u(z))}{c(w, x : y, z)} = 1 + O(t^c(u-1))$$

and

$$\frac{c(\tau_u(w), \tau_u(y) : \tau_u(x), \tau_u(z))}{c(w, y : x, z)} = u(1 + O(t^c(u-1))).$$

Proof. If w, x, y and z all lie in $v^{-1}(\mathbb{R}_{>0})$ then τ_u multiplies w, x, y and z by u, which doesn't change their cross-ratio. If w, x, y and z all lie in $v^{-1}(\mathbb{R}_{<0})$ then τ_u has no effect at all.

Suppose that $w \in v^{-1}(\mathbb{R}_{>0})$ and x, y and $z \in v^{-1}(\mathbb{R}_{<0})$. Then

$$\begin{aligned}\frac{c(\tau_u(w), \tau_u(x) : \tau_u(y), \tau_u(z))}{c(w, x : y, z)} &= \frac{(uw-y)(x-z)/(uw-z)(x-y)}{(w-y)(x-z)/(w-z)(x-y)}\\ &= \frac{(uw-y)(w-z)}{(uw-z)(w-y)}\\ &= c(uw, w : y, z)\end{aligned}$$

By Proposition 5.2.2, this is $1+O(t^d)$ where d is the length of the finite edge of the tree $[w, uw, y, z]$. The fact that y and $z \in v^{-1}(\mathbb{R}_{<0})$ and w and $uw \in v^{-1}(\mathbb{R}_{>0})$ forces this edge to contain $\overline{\mathcal{R}^2}$. The edge extends for a distance of c below $\overline{\mathcal{R}^2}$ and $v(u-1)$ above $\overline{\mathcal{R}^2}$. The case where $w \in v^{-1}(\mathbb{R}_{<0})$ and x, y and $z \in v^{-1}(\mathbb{R}_{>0})$ is practically identical.

Now suppose that w and x are in $v^{-1}(\mathbb{R}_{>0})$ and y and z are in $v^{-1}(\mathbb{R}_{<0})$.

$$\begin{aligned}\frac{c(\tau_u(w), \tau_u(x) : \tau_u(y), \tau_u(z))}{c(w, x : y, z)} &= \frac{c(uw, ux : y, z)}{c(w, x : y, z)}\\ &= \frac{c(uw, ux : y, z)}{c(w, ux : y, z)} \cdot \frac{c(w, ux : y, z)}{c(w, x : y, z)}\\ &= c(w, uw : y, z)c(x, ux : y, z)\end{aligned}$$

where the last equality is by the same identity as was used in the previous set of displayed equations. Each factor is $1 + (t^c(u-1))$ by the same argument as before.

Finally, we have

$$\begin{aligned}
\frac{c(\tau_u(w),\tau_u(y):\tau_u(x),\tau_u(z))}{c(w,y:x,z)} &= \frac{c(uw,y:ux,z)}{c(w,y:x,z)} \\
&= \frac{(uw-ux)(y-z)/(uw-z)(y-ux)}{(w-x)(y-z)/(w-z)(y-x)} \\
&= u\frac{(w-z)(y-x)}{(uw-z)(uy-x)}.
\end{aligned}$$

We have $\frac{w-z}{uw-z} = 1-(u-1)\frac{w}{uw-z}$. Since $v(z) < 0$ and $v(uw) = v(w) > 0$, we have $v(\frac{w}{uw-z}) = v(w/z) = v(w) - v(z) \geq c$, so $\frac{w-z}{uw-z} = 1 + O(t^c(u-1))$. By the same logic, $\frac{y-x}{uy-x} = 1 + O(t^c(u-1))$. □

5.8 Deformation of G and $Z_i^{\pm}$

Suppose that we are given a zero tension curve (Γ, w, ι) in $\mathbb{R}^n$. Our aim is to construct a curve $X \in K^n$ with $\operatorname{Trop} X = (\Gamma, \iota, w)$ by first finding a group G and subsets $Z_i^{\pm}$ as in the previous section and then deforming our choices until each $\mu_{Z_i^+, Z_i^-}$ is identically 1 as desired.

Let $\mathrm{Edge}(\Gamma^{\mathrm{fin}})$ denote the set of directed edges of Γ^{fin}. Fix a basepoint in Γ, let $\tilde{\Gamma}$ denote the universal cover of Γ with $\pi : \tilde{\Gamma} \to \Gamma$ the projection and $v_0 \in \tilde{\Gamma}$ a specified lift of the basepoint. So $\pi_1(\Gamma, \pi(v_0))$ acts on $\tilde{\Gamma}$.

Suppose that we have a map $h : \mathrm{Edge}(\tilde{\Gamma}^{\mathrm{fin}}) \to \mathrm{PGL}_2(K)$ which obeys $h(-e) = h(e)^{-1}$ and we have two finite vertices v and w of $\tilde{\Gamma}$, let $e_1, \ldots, e_r$ be the path between them. We define $p_h(v_1 \to v_2) = h(e_1) \cdots h(e_r)$. If h factors through the projection $\pi : \tilde{\Gamma} \to \Gamma$ then $g \mapsto p_h(v_0 \to gv_0)$ gives a map of groups

$\pi_1(\Gamma, \pi(v_0)) \to \mathrm{PGL}_2(K)$.

We now prove a proposition left open before. Remember our construction of the multisets $\tilde{Y}_i^{\pm}$: These are finite multisets each of whose elements is a degree 1 vertices of $\tilde{\Gamma}$. They have the property that, when the paths from $\tilde{y}_k^{+,i}$ to $y_k^{-,i}$ for are pushed down to Γ, the signed number of times these paths pass through the edge e in the direction running away from endpoint $v_1(e)$ is $\sigma_{v_1(e)}(e)_i$.

Proposition 5.8.1. *There is a Schottky group G and subsets $Z^{\pm,1}$, …, $Z^{\pm,1}$ of Ω such that, writing $Z = \bigcup_{i=1}^{n} Z^{\pm,i}$, the tree $\tilde{\beth}(G,Z)$ is isometric to $\tilde{\Gamma}$, $\beth(G,Z)$ is isometric to Γ and the ends of $\tilde{\beth}(G,Z)$ ending at the $Z_i^{\pm}$ correspond to the ends of $\tilde{\Gamma}$ ending at the $\tilde{Y}^{\pm,i}$.*

Proof. Let T be a tree with $j : T \to \Gamma$ a surjection preserving edge lengths which is bijective on edges (but some vertices may be multiply covered.) We may find an isometry $\psi : T \hookrightarrow \mathrm{BT}(K)$. Let $u_0 \in T$ be a preimage of $\pi(v_0)$ and orient each edge of T away from v_0. We write $v_1(e)$, $v_2(e)$ for the endpoints of e closer to and further from u_0 respectively. Proceeding inductively away from u_0, for each $e \in \mathrm{Edge}(T^{\mathrm{fin}})$ we may choose a map $h(e)$ such that $p_h(v_0 \to v)\psi(v_0) = \psi(v)$ for all $v \in T$. More explicitly, when we come to each edge e, we must make sure that $h(e)(p_h(v_0 \to v_1)^{-1}(\psi(v_1))) = p_h(v_0 \to v_1)^{-1}(\psi(v_2))$. We require furthermore that $h(e)$ is hyperbolic and that the path $[V_1(h(e)), V_2(h(e))]$ meet T along the edge $[(p_h(v_0 \to v_1)^{-1}(\psi(v_2))), (p_h(v_0 \to v_1)^{-1}(\psi(v_1)))]$. If we make our choices generically, the following additional condition will also hold:

> If $j(v) = j(v')$ and e and e' are distinct edges directed away from v and v' then $[V_1(h(e)), V_2(h(e))] \cap [V_1(h(e')), V_2(h(e'))] = \psi(v_0)$.

We now extend h to all of $\mathrm{Edge}(\tilde{\Gamma}^{\mathrm{fin}})$ by requiring that h factor through π. So we get a map of groups $\pi_1(\Gamma, \pi(v_0)) \to \mathrm{PGL}_2(K)$. We also define a map $\psi_h : \mathrm{Vert}(\tilde{\Gamma}) \to \mathrm{BT}(K)$ by $\psi_h(v) = p_h(v_0 \to v)\psi(v_0)$; by the above, this extends ψ. The map ψ_h is obviously compatible with the $\pi_1(\Gamma, \pi(v_0))$ action. Using this compatibly to translate back to T, we see that ψ_h can be extended to $\tilde{\Gamma}$.

We claim that this extension is still injective and is an isometry onto its image. ψ_h is locally injective, as it must be locally injective in the neighborhood of a point in the middle of an edge and our "additional condition" guarantees that it is injective in the neighborhood of a vertex. That means that a non-backtracking path in $\tilde{\Gamma}$ is taken to a non-backtracking path in $\mathrm{BT}(K)$. Nonbacktracking paths are geodesics in trees like $\tilde{\Gamma}$ and in $\mathbb{R}$-trees like $\mathrm{BT}(K)$, so ψ_h takes geodesics to geodesics. Since geodesics are unique in $\mathrm{BT}(K)$, the map ψ_h is not merely length preserving on edges but distance preserving on all of $\tilde{\Gamma}$. Isometries are injective.

This shows that $p_h(\pi_1(\Gamma, \pi(v_0)))$ is free: it acts freely on the tree $\psi_h(\tilde{\Gamma})$ and it is Schottky as it is easy to check that any nonhyperbolic element, when acting on a tree containing its fixed point(s), fixes a line segment. Also, ψ_h is an isometry because ψ was. $\square$

Now, let $u : \mathrm{Edge}(\Gamma^{\mathrm{fin}}) \to U := \{x \in K^* : v(x) = 0\}$. We will replace h by a new function $h^u : \mathrm{Edge}(\Gamma^{\mathrm{fin}}) \to \mathrm{PGL}_2(K)$ defined as follows: $h^u(e)$ has the same eigenvectors as $h(e)$ but $\lambda_1(h^u(e))/\lambda_2(h^u(e)) = u\lambda_1(h(e))/\lambda_2(h(e))$. We define a new map $\psi^u : \mathrm{Vert}T \hookrightarrow \mathrm{BT}(K)$ by $v \mapsto p_{h^u}(v_0 \to v)\psi(v_0)$.

Lemma 5.8.2. *The map ψ^u extendeds to an isometry $T \hookrightarrow \mathrm{BT}(K)$. h^u then gives a set of choices compatible with ψ^u in the construction of the previous theorem.*

Proof. To check this, one first needs to see that, for v_1 and v_2 endpoints of an edge e, the distance between $\psi^u(v_1)$ and $\psi^u(v_2)$ is the length of e. Using the isometry $p_{h^u}(v_0 \to v_1)^{-1}$, this is transformed into a problem about computing the distance from $\psi(v_0)$ to $h^u(e)\psi(v_0)$. Since the path between the eigenvectors of $h^u(e)$, which are also the eigenvectors of $h(e)$, passes through $\psi(v_0)$, this distance is $v(\lambda_1(h^u(e))/\lambda_2(h^u(e))) = v(\lambda_1(h(e))/\lambda_2(h(e)))$ as $v(u) = 0$. Reversing the argument, the latter distance is that from $\psi(v_1)$ to $\psi(v_2)$, which is already assumed to be correct.

One must also check that, for each vertex v of T, the paths joining $\psi^u(v)$ to the images of the neighbors of v meet only at v. Using $\psi^u(v_0 \to v)^{-1}$ to translate the problem back to v_0, we need to check that the arcs connecting $\psi(v_0)$ to $V_2(h^u(e))$ meet only at $\psi(v_0)$ as e varies. But the eigenvectors of $h^u(e)$ are the same as those of $h(e)$, so again this reduces to a condition that has already been checked for ψ.

The last sentence is straightforward. $\square$

We may then define the embedding $\psi^u : \tilde{\Gamma} \to \mathrm{BT}(K)$ as before. When we change h to h^u, we also change the $Z^{\pm}$ to be at the corresponding ends of $\psi^u(\tilde{\Gamma})$.

Fix $\gamma_0 \in \pi_1(\Gamma, \pi(v_0))$. Our goal is to determine the effect of all of these changes on $\mu_{Z_i^+, Z_i^-}(\gamma_0)$, which is given by the formula

$$\mu_{Z_i^+, Z_i^-}(\gamma_0) = \prod_{k=1}^{d_i} \prod_{\gamma \in \pi_1(\Gamma, \pi(v_0))} c(w, \gamma_0(w) : \gamma z_i^{+,k}, \gamma z_i^{-,k})$$

for an arbitrary choice of w.

Since w may be taken arbitrarily, we choose a w such that the paths from w to both $\psi(\tilde{\Gamma})$ and $\psi^u(\tilde{\Gamma})$ attach at v_0. Let $\mathcal{T}$ be the tree which consists of $\tilde{\Gamma}$ with two more infinite edges attached at v_0 and gv_0. We denote the two additional endpoints by a and b respectively. We extend the embeddings ψ and ψ^u to take $\mathcal{T} \hookrightarrow \mathrm{BT}(K)$ by taking the infinite edge attached at v_0 to w and the infinite edge attached at gv_0 to $p_h(g)w$ and $p_{h^u}(g)w$ respectively.

The main computation of this section is the following:

Proposition 5.8.3. *Let c be the length of the shortest edge of Γ^{fin}. For e an edge of Γ with endpoints $v_1(e)$ and $v_2(e)$, let $[\gamma_0, e]$ be the number of times a closed loop representing γ_0 passes through e in the direction away from $v_1(e)$. Then replacing h by h_u changes $\mu_{Z^+,Z^-}(\gamma_0)$ by a factor of*

$$\prod_{e \in \Gamma^{\mathrm{fin}}} u(e)^{[\gamma_0, e]\sigma_{v_1(e)}(e)_i} (1 + O(t^c \min_{e \in \Gamma^{\mathrm{fin}}} (u(e) - 1)))$$

Proof. We will describe a sequence ψ^j of embeddings $\mathcal{T} \hookrightarrow \mathrm{BT}(K)$ with $\psi^0 = \psi$ and such that, for every point c of $\mathcal{T}$, we have $\psi^j(c) = \psi^u(c)$ for j sufficiently large. More precisely, define an ordering $e_1, e_2, \ldots$ on the edges of $\tilde{\Gamma}^{\text{fin}}$ such that each initial segment is a subtree of $\mathcal{T}$ containing v_0. Define $h^j(e_i)$ for $e \in \tilde{\Gamma}^{\text{fin}}$ to be $h^u(e_i)$ when $i \leq j$ and $h(e_i)$ otherwise. h^j does not factor through π, so we do not get a map $\pi_1(\Gamma, \pi(v_0)) \to \mathrm{PGL}_2(K)$, but we still do get an injection $\mathrm{Vert}(\tilde{\Gamma}^{\text{fin}}) \hookrightarrow \mathrm{BT}(K)$ by $v \mapsto p_h(v_0 \to v)\psi(v_0)$ which still extends to a map $\psi^i : \mathcal{T} \hookrightarrow \mathrm{BT}(K)$. The only subtlety is what to do with the infinite edges. For each infinite edge e of $\mathcal{T}$, let v be the vertex at which it is attached and send e to $p_{h^j}(v_0 \to v)p_h(v_0 \to v)^{-1}e$. Clearly, for every point c of $\mathcal{T}$, we have $\psi^j(c) = \psi^u(c)$ for j sufficiently large.

The map ψ^j is obtained from ψ^{j-1} as follows: let $f \in \mathrm{PGL}_2(K)$ be a map that takes $\psi^{j-1}(v_2(e))$ to $\overline{\mathcal{R}^2}$ and takes the eigenvectors of $p_{h^j}(v_0 \to v_2(e_j))p_{h^j}(v_0 \to v_1(e_j))^{-1}$ to 0 and ∞. Then $\psi^j = f^{-1}\tau_u f \psi^{j-1}$ where τ_u is the operator from Lemma 5.7.2.

Recall that $\tilde{Y}_i^{\pm}$ is the multiset of degree 1 vertices of $\mathcal{T}$ which ψ takes to $Z_i^{\pm}$. Set

$$\mu^j = \prod_{k=1}^{d_i} \prod_{\gamma \in \pi_1(\Gamma, \pi(v_0))} c(\psi^j(a), \psi^j(b) : \psi^j(\gamma \tilde{y}_i^{+,k}), \psi^j(\gamma \tilde{y}_i^{-,k})).$$

For all but finitely many γ, the subtree of $\mathcal{T}$ spanned by a, b, $\gamma \tilde{y}_i^{+,k}$ and $\gamma \tilde{y}_i^{-,k}$ has an internal edge dividing a and b from $\gamma \tilde{y}_i^{+,k}$ and $\gamma \tilde{y}_i^{-,k}$. Then that $c(\psi^j(a), \psi^j(b) : \psi^j(\gamma \tilde{y}_i^{+,k}), \psi^j(\gamma y_i^{-,k})) = 1 + O(t^d)$ where d is the length of that edge. Moreover, for any fixed $D \in \mathbb{R}$, there are only finitely many terms for which $d < D$. Thus, we see that the product is uniformly convergent. (In the nonarchimedean topology, any product where the terms approach 1 converges.) Thus, we may take the limit as $j \to \infty$ term by term and get

$$\lim_{j\to\infty} \mu^j = \prod_{k=1}^{d_i} \prod_{\gamma \in \pi_1(\Gamma, \pi(v_0))} c(\psi^u(a), \psi^u(b) : \psi^u(\gamma \tilde{y}_i^{+,k}), \psi^u(\gamma \tilde{y}_i^{-,k})).$$

Our goal is to estimate the ratio of the above product to μ^0. We first estimate μ^j/μ^{j-1}. We claim that

$$\mu^j/\mu^{j-1} = u(\pi(e_j))^{\sigma_{v_1(e_j)}(e_j)_i}(1 + O(t^c(u(\pi(e_j)) - 1)))$$

if e_j is on the path from a to b and

$$\mu^j/\mu^{j-1} = 1 + O(t^c(u(\pi(e_j)) - 1))$$

otherwise.

This is a product of cross ratios as in Lemma 5.7.2. Every term where the tree spanned by a, b, $\gamma\tilde{y}_i^{+,k}$ and $\gamma\tilde{y}_i^{-,k}$ does not contain the edge e_j is 1. Every term where e_j does not separate a from b and $\gamma\tilde{y}_i^{+,k}$ from $\gamma\tilde{y}_i^{-,k}$ is $1 + O(t^c(u(e_j) - 1))$. The bound of the valuations by c is because, all the terms that are on the $v_1(e_j)$ side of e_j are at least the length of e_j, which is $\geq c$, away from the $v_2(e)$ side. So, if e_j does not separate a from b, we have the claimed result.

Suppose now that e_j does seperate a from b. For the accuracy that concerns us, the only terms we need to consider are the ones coming from when e_j also seperates $\gamma\tilde{y}_i^{+,k}$ from $\gamma\tilde{y}_i^{-,k}$. For each such term, we get a contribution of $u(e_j)(1 + O(t^c(u(e_j) - 1)))$. The number of γ for which this happens (counted with sign) is $\sigma_{v_1(e_j)}(e_j)_i$. So we have proven the claimed bound for μ^j/μ^{j-1}.

We want to estimate $\prod_{j=1}^{\infty} \mu^j/\mu^{j-1}$. The terms for which e_j does not seperate a from b only contribute to the error. For a given edge $e \in \Gamma^{\mathrm{fin}}$, the number of e_j for which $\pi(e_j) = e$ and e_j seperates a from b is $[\gamma_0, e]$. So we get the claimed result. □

If we now put together all of our results, we will have proven Theorem 5.0.4 for curves of genus $g \geq 2$. Let us lay out the argument in summary. Proposition 5.6.4 states that, if we can find G and $Z^{\pm,1}$, ..., $Z^{\pm,n}$ such that $\tilde{\beth}(G,Z)$ is isometric to $\tilde{\Gamma}$ and such that each $\mu_{Z^{+,i},Z^{-,i}}$ is identically zero, then the image of the map $\phi : \Omega \to (K^*)^n$ is a genus g curve with tropicalization $\iota(\Gamma)$. Proposition 5.8.1 tells us that we can find G and $Z^{\pm,i}$ such that $\tilde{\beth}(G,Z)$ is isometric to $Ga\tilde{m}ma$.

In this section, we have described a way of making $U^{\mathrm{Edge}(\Gamma)}$ act on the collection of possible choices of G and $Z^{\pm,i}$ while preserving the geometry of $\tilde{\beth}(G,Z)$. Choose a generating set γ_1, ..., γ_g for $\pi_(\Gamma, \pi(v_0))$. If $\mu_{Z^{+,i},Z^{-,i}}(\gamma_j)$ is 1 for $1 \leq j \leq g$ then we will know that $\mu_{Z^{+,i},Z^{-,i}}(\gamma)$ is identically 1. Proposition 5.8.3 showed that the action of $U^{\mathrm{Edge}(\Gamma)}$ modifies $\mu_{Z^{+,i},Z^{-,i}}(\gamma_j)$ by a factor of

$$\prod_{e \in \Gamma^{\mathrm{fin}}} u(e)^{[\gamma_j,e]\sigma_{v_1(e)}(e)_i}(1 + O(t^c \min_{e \in \Gamma^{\mathrm{fin}}}(u(e)-1)))$$

where c was the length of the shortest edge of Γ. This is an action on U^{gn} of the form discussed in Proposition 5.7.1. The matrix A in this case is precisely the matrix $\mathrm{Slope}(\Gamma, \iota, wt)$.

Our hypothesis that Γ is ordinary tells us that $A = \Phi$ has full rank and we have imposed that it has full rank modulo p as well. So Proposition 5.7.1 tells us that the action of U^E on U^{gn} is transitive. In particular, it is possible to choose a u which makes all of the $\mu_{Z^{+,i},Z^{-,i}}(\gamma_j)$ equal to 1. We thus have proven Theorem 5.0.4.

Bibliography

[1] F. Ardila and C. Klivans: "The Bergman complex of a matroid and phylogenetic trees" `arXiv:math.CO/0311370`

[2] D. Bayer and I. Morrison: "Standard Bases and Geometric Invariant Theory I: Initial Ideals and State Polytopes" *Journal of Symbolic Computation* **6** (1988) 209–217

[3] G. Bergman: "The logarithmic limit-set of an algebraic variety", *Transactions of the American Mathematical Society* **157** (1971) 459–469.

[4] V. Berkovich: *Spectral Theory and Analytic Geometry over non-Archmedean Fields* Mathematical Surveys and Monographs v. 33 Amer. Math. Soc. 1990

[5] J.-P. Barthélemy and A. Guénoche: *Trees and Proximity Representations*, Wiley-Interscience, Chichester, 1991.

[6] R. Bieri and J.R.J. Groves: "The geometry of the set of characters induced by valuations" *J. reine und angewandte Mathematik* **347** (1984) 168–195.

[7] L.J. Billera, S. Holmes and K. Vogtmann: "Geometry of the space of phylogenetic trees" *Advances in Applied Mathematics* **27** (2001) 733–767.

[8] T. Brylawski: "A Combinatorial Model for Series-Parallel Networks" *Transactions of the AMS* **154** (1971) 1–22

[9] P. Buneman: "A note on metric properties of trees" *Journal of Combinatorial Theory, Ser. B* **17** (1974) 48–50.

[10] H. Crapo: "A Higher Invariant for Matroids" *Journal of Combinatorial Theory* **2** (1967) 406–417

[11] M. Develin and B. Sturmfels: "Tropical Convexity" *Documenta Math.* **9** (2004) 1–27

[12] A. Dress: "Duality Theory for Finite and Infinite Matroids with Coefficients" *Advances in Mathematics* **59** (1986) 97–123

[13] A. Dress and W. Wenzel: "Valuated matroids" *Advances in Mathematics* **93** (1992) no. 2 214–250

[14] A. Dress and W. Wenzel: "Grassmann-Plücker Relations and Matroids with Coefficients" *Advances in Mathematics* **86** (1991) 68–110

[15] A. Dress and W. Terhale: "The Tree of Life and Other Affine Buildings", *Documenta Mathematica* 1998 Extra Vol. III, 565–574 (electronic)

[16] L. Ein: “Hilbert scheme of smooth space curves” *Ann. Sci. École Norm. Sup. (4)* **19** (1986) no. 4 469–478

[17] M. Einsiedler, M. Kapranov and D. Lind: “Non-archimedean amoebas and tropical varieties”, `arXiv:math.AG/0408311` preprint, submitted to Crelle’s Journal

[18] D. Eisenbud: *Commutative Algebra with a View Toward Algebraic Geometry* New York: Springer-Verlag 1995

[19] Gerritzen, L. *Schottky groups and Mumford curves* New York: Springer-Verlag 1980

[20] I. Gelfand, I. Goresky, R. MacPherson and V. Serganova: “Combinatorial geometries, convex polyhedra, and Schubert cells”, *Adv. Math.* **63** (1987) 301–316

[21] I. Gelfand, M. Kapranov and A. Zelevinsky: *Discriminants, Resultants and Multidimensional Determinants* Boston: Birkhäuser 1994

[22] M. Kalkbrener and B. Sturmfels: “Initial Complexes of Prime Ideals” *Advances in Mathematics* **116** (1995) no.2 365–376

[23] D. Maclagan, “Antichains of Monomial Ideals are Finite” *Proceedings of the American Mathematical Society* **129** (2001) 1609–1615

[24] G. Mikhalkin: “Counting curves via lattice paths in polygons” *C. R. Acad. Sci. Paris* **336** (2003) no. 8 629–634

[25] G. Mikhalkin: “Enumerative Tropical Algebraic Geometry” Preprint 2005

[26] Y. Manin and V. Drinfeld: “Periods of p-adic Schottky groups” *J. Reine Angew. Math.* **262/263** (1973) 239–247

[27] T. Mora and L. Robbiano: “The Gröbner Fan of an Ideal” *Journal of Symbolic Computaion* **6** (1998) 183–208

[28] J. Morgan and P. Shalen: “Valuations, trees, and degenerations of hyperbolic structures I” *Ann. of Math. (2)* **120** (1984) no. 3 401–476

[29] D. Mumford: “An analytic construction of degenerating curves over complete local rings” *Compositio Math.* **24** (1972) 129 – 174

[30] T. Nishinou and B. Siebert: “Toric degenerations of toric varieties and tropical curves” `arXiv:math.AG/0409060`

[31] L. Pachter and D. Speyer: “Reconstructing Trees from Subtree Weights” *Applied Mathematics Letters* **17** (2004) 615 – 621.

[32] M. Raynaud and L. Gruson: “Critères de Platitude et de Projectivité” *Invent. Math.* **13** 1–89 (1971)

[33] A. Robinson and S. Whitehouse: “The tree representation of Σ_{n+1}” *Journal of Pure and Applied Algebra* **111** (1996) 245–253

[34] J. Silverman: *Advanced topics in the theory of elliptic curves* Grad. Texts in Math. vol. 151, Springer-Verlag: New York 1994

[35] A. Smirnov: "Torus schemes over a discrete valuation ring" *St. Petersburg Math. J.* **8** (1997) no. 4 651–659 (English translation of Russian original, see *Algebra i Analiz* **8** (1996) no. 4 161–172 for original paper.)

[36] F. Sottile: "Enumerative real algebraic geometry. Algorithmic and quantitative real algebraic geometry" 139–179, DIMACS Ser. Discrete Math. Theoret. Comput. Sci., 60, Amer. Math. Soc.: Providence 2003

[37] D. Speyer: "Horn's Problem, Vinnikov Curves and the Hive Cone", *Duke Mathematical Journal*, to appear. `arXiv:math.AG/0311428`

[38] D. Speyer and B. Sturmfels: "Tropical Mathematics" `arXiv:math.CO/0408099`

[39] D. Speyer and B. Sturmfels: "The Tropical Grassmannian", *Advances in Geometry* **4** (2004) 389-411

[40] B. Sturmfels: *Algorithms in Invariant Theory*, Springer: Vienna, 1993

[41] B. Sturmfels: *Gröbner Bases and Convex Polytopes*, University Lecture Series **8**, American Mathematical Society, 1996

[42] B. Sturmfels: *Solving Systems of Polynomial Equations*, CBMS Regional Conference Series in Math., **97**, American Mathematical Society, 2002

[43] B. Sturmfels: "Gröbner Bases of Toric Varieties" *Tohoku Math. J.* **43** (1991) no. 2 249–261

[44] B. Sturmfels: "Viro's theorem for complete intersections", *Annali della Scuola Normale Superiore di Pisa (4)* **21** (1994) no. 3 377–386

[45] J. Tevelev: "Tropical Compactifications" `arXiv:math.AG/0412329`

[46] O. Viro: "Dequantization of Real Algebraic Geometry on a Logarithmic Paper", Proceedings of the 3rd European Congress of Mathematicians, Birkhüser, Progress in Math, **201** (2001) 135–146

[47] O. Viro: *Patchworking Real Algebraic Varieties*, Preprint Uppsala University U.U.D.M. Report 1994:42. Also available at `http://www.math.uu.se/~ oleg`

[48] K. Vogtmann: "Local structure of some OUT(F_n)-complexes" *Proceedings of the Edinburgh Mathematical Society* **33** (1990) 367–379

[49] N. White: *Theory of Matroids – Encyclopedia of Mathematics and its Applications vol. 26* Cambridge University Press: London (1986)

[50] N. White: *Matroid Applications – Encyclopedia of Mathematics and its Applications vol. 40* Cambridge University Press: London (1992)

[51] G. Ziegler: *Lectures on Polytopes* Graduate Texts in Mathematics vol. 152, Springer-Verlag: New York 1995